TEC-0127

Site Model Based Image Registration and Change Detection

R. Chellappa P. Burlina
C.L. Lin X. Zhang
L.S. Davis A. Rosenfeld

University of Maryland
Center for Automation Research
College Park, MD 20742-3275

November 1999

Approved for public release; distribution is unlimited.

Prepared for:
Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714

Monitored by:
U.S. Army Corps of Engineers
Topographic Engineering Center
7701 Telegraph Road
Alexandria, VA 22315-3864

US Army Corps of Engineers
Topographic Engineering Center

T
E
C

Destroy this report when no longer needed.
Do not return it to the originator.

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

The citation in this report of trade names of commercially available products does not constitute official endorsement or approval of the use of such products.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE November 1999	**2. REPORT TYPE** Technical, Final	**3. DATES COVERED** *(From - To)* September 1992 - March 1997

4. TITLE AND SUBTITLE

Site Model Based Image Registration and Change Detection

5a. CONTRACT NUMBER
DACA76-92-C-0024

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)

R. Chellappa, P. Burlina, C.L. Lin, X. Zhang, L.S. Davis, A. Rosenfeld

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Maryland
Center for Automation
Research
College Park, MD 20742-3275

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research
Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714

U.S. Army Topographic
Engineering Center
7701 Telegraph Road
Alexandria, VA 22315-3864

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
TEC-0127

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The model-based detection schemes presented in this report incorporate Image Understanding (IU) techniques that can be directed by the image analyst to conduct spatially-constrained analyses, whose outcomes may be indicative of occurrences of change with special intelligence significance. The aerial image exploitation system is site model driven, and is generally based on three classes of primitives: object primitives, spatial primitives, temporal primitives. This work emphasizes the use of geometric (i.e., object and spatial) primitives. These take the form of site models implemented in the RADIUS Common Development Environment (RCDE). The models encode the spatial relationships between fixed objects of interest in a site, such as buildings, roads, etc.

15. SUBJECT TERMS

Change detection, site monitoring, vehicle counting, image positioning

16. SECURITY CLASSIFICATION OF: UNCLASSIFIED			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Jean Diaz
a. REPORT UNCLASSIFIED	**b. ABSTRACT** UNCLASSIFIED	**c. THIS PAGE** UNCLASSIFIED	UNLIMITED	56	**19b. TELEPHONE NUMBER** *(include area code)* (703) 428-6022

TABLE OF CONTENTS

List of Figures

List of Tables

Preface

This research was sponsored by the Defense Advanced Research Projects Agency (DARPA Order No. 8979) and monitored by the U.S. Army Topographic Engineering Center (TEC) under contract DACA 76-92-C-0024, titled "Site Model Based Image Registration and Change Detection". The ARPA Program Manager was Dr. Thomas Strat, and the TEC Contracting Officer's Representative was Ms. Lauretta Williams.

Acknowledgements

We wish to acknowledge Aaron Heller of SRI for helping us with the RCDE, and Richard Welty of GE for providing invaluable help in the integration of the algorithm.

SITE MODEL BASED IMAGE REGISTRATION AND CHANGE DETECTION

1. INTRODUCTION

The overriding concern of RADIUS research has been the development of algorithms for model-based and context-based exploitation of aerial imagery. The context-based exploitation paradigm has proven to be important for the design of less complex and more robust change detection and monitoring tools. Change detection (CD) and monitoring, with the goals of locating and identifying significant changes or relevant activities that have occurred between the times of acquisition of the imagery, are core aerial image analysis operations (Strat and Climenson, 1994). Previous efforts (Kawamura, 1971; Price, 1982; Ulstad, 1973) have relied on general purpose methods to screen a variety of imagery without access to any site-specific information. Methods such as these have proven to be unreliable because too many inconsequential changes occur in any natural environment; therefore, monitoring techniques based on more or less sophisticated and sometimes low-level pixel-based methods are extremely sensitive to registration errors and photometric conditions. Even if general-purpose methods could be developed for screening out all changes due to variations in viewpoint and illumination, many differences between the images would still be present whose significance could only be determined by an image analyst (IA) with comprehensive site knowledge.

The model-based detection schemes presented in this report incorporate Image Understanding (IU) techniques that can be directed by the IA to conduct spatially-constrained analyses, whose outcomes may be indicative of occurrences of change with special intelligence significance. The aerial image exploitation system is site model driven, and is generally based on three classes of primitives: *object primitives*, which correspond to the specific objects that occur in a particular site model and to the generic object classes supported by the IU system; *spatial primitives*, for the construction of search locales and the specification of constraints on the search for object types within locales; and *temporal primitives*, which can constrain or parameterize the analysis by factors such as time of day, day of week, time of year, etc. This work emphasizes the use of geometric (i.e., object and spatial) primitives. These take the form of site models implemented in the RADIUS Common Development Environment (RCDE) (SRI, 1993). The models encode the spatial relationships between fixed objects of interest in a site, such as buildings, roads, etc.

IU algorithms, designed to extract objects such as buildings or vehicles in a site for CD applications, cannot be purely bottom up. Site models are instrumental in enabling the incorporation of feedback mechanisms into IU algorithms
(Bejanin *et al.*, 1994; Collins *et al.*, 1994; McKeown and Roux, 1994). For example, in extracting buildings (Venkateswar and Chellappa, 1992), heuristics based on the expected shapes of roofs (site-specific information) are employed for completing any partial roof hypotheses that result from imperfect bottom-up processing. Likewise, shadow analysis is important for obtaining height information

(Huertas and Nevatia, 1988; McKeown, 1990), or allowing the IU system to explain why some building features that are in the field of view cannot be identified in the image. Similarly, site models can provide geometric and photometric constraints that reduce matching ambiguities or search spaces; therefore, the use of site model information increases the reliability and decreases the complexity of IU processes.

Among the problems addressed in the RADIUS research, three main areas of interest have been singled out as critical for the project: construction of site models, automated positioning of images, and monitoring of movable objects. The focus of our group has been on monitoring movable objects. More specifically, we have addressed the following two tasks: (a) detecting and counting vehicles in selected locales, and (b) monitoring and characterizting vehicle groupings.

In contrast to earlier CD techniques (Kawamura, 1971; Price, 1982; Ulstad, 1973), we are using model-based detection schemes
(Bejanin *et al.*, 1994; Collins *et al.*, 1994; Huertas and Nevatia, 1988)
(McKeown and Roux, 1994; McKeown, 1990) that exploit known geometric models. Different approaches have been employed based on the task at hand: An analytical recognition scheme is used for task (a) and a global method using spectral analysis is employed for task (b). Site model information is incorporated in different ways and to various extents according to the selected detection strategy. The emphasis in what follows is on providing a general description of our system, its components, and integration issues.

The vehicle detection and counting tool was originally developed to illustrate the Model Supported Exploitation (MSE) approach and the use of locale and object model information. Vehicles have simple box-like 3-D models with specified dimensions. A modified Canny edge detector (Venkateswar and Chellappa, 1992) is used to detect edges and gradient directions. A modified generalized Hough transform is used to locate candidate vehicles. For each candidate rectangle obtained, a rubberband template is generated by the 3-D vehicle model and compared with a local edge map. Lastly, a shadow verification method is used to confirm the detection of vehicles. Three-dimensional object model and site information (camera model, illuminant, etc.) is exploited throughout the detection and counting procedure.

The module for detecting vehicle configuration illustrates the use of spectral information derived from spatial models. The purpose of this module is principally to serve as an attentional mechanism for detecting object configurations of particular significance. This operator is applied to detecting the presence of convoys on roads, occupancy of parking areas, regular vehicle formations in staging areas, and trains on rails. The problem of detecting geometrical configurations has been reported in papers such as (Bronskill *et al.*, 1989; Khang and Robins, 1991; Xu, 1988), using knowledge-based association or Hough transform schemes. Our method uses a frequency domain approach. We look at the spectral representation associated with directional edges (in dominant directions over an area) to determine any periodic behavior of this underlying structure. This method uses available site model information and enforces constraints derived from object model geometry in the frequency domain. This operator is used stand alone or as a pre-processor to the local detection and counting module.

An important part of our work has been assessing the sensitivity of the detection modules to tuning parameters and site model information. Since these algorithms are to be used by IAs in a batch mode over large sets of images, we have provided means to incorporate automatic tuning and training mechanisms into our modules, thereby avoiding manual parameter tuning as much as possible.

Another important facet of our work has been system integration within the RCDE. One of the principal thrusts behind this research has been the development of IU capabilities focused on the needs of IAs. To this end, the mechanisms previously described have been developed within or around the RCDE platform, making heavy use of its integrated functionalities; this platform provides a common development environment on which RADIUS-related algorithms can be developed and tested, and allows for the creation and manipulation of CAD-like objects used to model site objects. It also provides a framework from which the RADIUS testbed system is derived. A typical exploitation cycle in this platform involves the following steps: new images are acquired, and they are registered to the existing site model. Images then are prioritized for further analysis. The IA launches a set of batch detection algorithms; the user specifies from the site model which areas of interest should be considered for exploitation. The needed locale and object information is fetched from site objects residing in a database. The successive steps involved in the exploitation cycle — registration, resection, enhancement, and detection — are implemented in or around RCDE. Integration into RCDE has enabled successful porting of site data and algorithms.

This report is organized as follows: Section 2. gives a detailed overview of site-model-based CD as well as the specific type of site model information used by each CD process described subsequently. The model-based approach to detection and counting of vehicles is reported in Section 3. The detection of convoys and vehicle formations is reported in Section 4. Experimental results using real imagery are reported in each of these sections. Sensitivity analysis and the statistical characterization of the algorithms' performance, as well as issues related to parameter training, are presented in Section 5. Finally, issues dealing with integration are reported in Section 6..

2. SITE MODEL SUPPORTED MONITORING

A core component of MSE is the availability of site models. A site model contains a geometrical description of the site under scrutiny and of relevant site features (areas, buildings, roads, etc.); it also includes imaging and photometric parameters associated with the available images along with collateral and auxiliary information (ARPA, 1993). Typical auxiliary information associated with a site model includes: (a) an overview of the site, (b) a baseline description, and (c) comments and analysis tasks. Site model construction requires several overlapping coverage images of the same site; it is carried out under RCDE.

When a newly acquired image becomes available, its registration to the existing site model is a necessary condition for model-based exploitation. As mentioned earlier, depending on the particular exploitation task, e.g., if building- or vehicle-related activities are monitored, we can use the site model and viewing direction of the new image to identify regions in the image that need further analysis. We subsequently can invoke the necessary IU algorithms for detecting and monitoring construction activities, as well as for locating and counting vehicles (and road extraction, if construction of roads is being monitored).

Perhaps one of the most obvious ways in which site models are used is in the delineation of Regions of Interest (ROIs). Given an image to be exploited, locating the regions of interest according to the task narrows the search area and reduces computational complexity and false alarms. Delineation not only decreases complexity, but also ensures the success of algorithms which otherwise might fail. Finally, using context in delineation allows for recognition by function, since location is a determinant factor in the ultimate purpose or function of an object. This is especially critical in applications where target signatures do not allow for easy object discrimination, such as in SAR images. The delineation process is straightforward in nature: when a 3-D region object is available from the site model, we directly project the region boundaries onto the image to label the region(s) in the image domain. The method uses camera model information available from the site model. The delineated region is then further cropped for possible shadows and/or occlusions, again exploiting site model information.

In addition to using site and camera information to delineate ROIs, the bulk of the site model information used by each detection algorithm is geometric in nature. Vehicle dimensions and orientation are used to infer template dimensions and orientation in the vehicle detection scheme. Selection of templates can be decided from a knowledge of the parking lot occupancy (in a full parking lot, the vehicle sides are likely to be occluded). This can be in the form of priors using context information, or this information can be fed back from the convoy detector. Information on climatic/weather conditions is not exploited here but evidently could explain the strengths of brightness gradients; this could be used either for tuning purposes or to quantify the confidence in the detected changes. The availability of CAD-like models for the visible objects allows for the application of syntactic object recognition techniques. Illumination direction is used along with camera parameters and object descriptions in the construction monitoring module to check for the presence of shadows corroborating the detected vertical lines. Contextual information also is used in the RCDE system to guide the application of exploitation algorithms. Conditions for which an algorithm's performance is acceptable can be encoded as rules embedded in a production system encoded in Prolog (Strat, 1995). Context information then is used to trigger the application of the algorithm deemed most appropriate for the task at hand.

3. VEHICLE DETECTION AND COUNTING

The vehicle detector can be triggered or cued by the convoy detector (see Section 4.) or can be used independently to monitor areas such as parking lots, roads, or training grounds. The general vehicle detection scheme relies on contour matching using information derived from the geometric model. The procedures used to carry out vehicle detection are reported in this section.

Since, in our implementation, we are primarily concerned with high-altitude imagery, vehicles are modeled as 3-D box-like objects with given width, length, and height specifications. The implementation of our vehicle detector consists of a pre-screener, an extractor and a verifier. Throughout the detection procedure, 3-D object models and site information (camera model, illuminant, etc.) are used. Details of the three stages of the vehicle detection process are now described.

3.1 Edge Detection and Prescreener

Edge detection is implemented by a modified version of the Canny edge detector. In particular, computations are limited to ROIs where the detection is carried out.

A modified generalized Hough transform (GHT) is used to locate areas corresponding to the center of candidate vehicles. The goal is to have edge pixels vote for possible loci of reference points. In our case, the reference point of an edge pixel is the center of a vehicle contour that contains the pixel. The displacement vector from any edge pixel to its reference point is represented using polar coordinates, indexed by the gradient direction of the edge pixel. The pre-screening algorithm can be described as follows: edge pixels are first extracted using the Canny edge detector. Both gradient magnitude and gradient direction are computed. A Hough table then is constructed for all vehicle templates. An accumulator array of possible reference points is created, and initialized to zero. For an edge pixel (x, y), we use its gradient direction ϕ to retrieve its associated displacement vectors from the Hough table and derive the positions of all possible reference points; thus, each edge pixel casts its vote for all reference points (x_c, y_c), where

$$\begin{aligned} x_c &= x + r(\phi)\cos[\alpha(\phi)], \\ y_c &= y + r(\phi)\sin[\alpha(\phi)]. \end{aligned}$$

The result of this step is a set of hypothesized vehicle centers, $\mathcal{L}$. Each point in the set $\mathcal{L}$ represents the center of a possible vehicle that has a match ratio (between its boundary contour and edge map) above a predetermined conservative threshold. Since we adopt a threshold scheme instead of searching for peaks in Hough space, we proceed with the detailed examination of vehicle contours.

3.2 Extractor

We subsequently apply local "rubberband" contour matching. This stage consists of four steps: (1) determining the positions and orientations of 3-D vehicle models; (2)

computing contours of the 3-D models; (3) computing textural features, and combining the features into a discrimination statistic that measures how "target-like" the detected object is; and (4) clustering into geometrically consistent objects. The algorithm first derives the positions and orientations of 3-D models. Centers of 3-D models need to be consistent with the back-projections of points in $\mathcal{L}$. Their orientations also are constrained to lie along the dominant direction of the active area.

For matching purposes, rubberband matching is used instead of tolerance band matching. The major disadvantage of tolerance-band matching is that it repetitively counts all pixels within tolerance bands. On the other hand, rubberband contours are deformable and will not overrate pixels clustered along the directions normal to the template contours. In addition, the degree of matching is evaluated separately on the vehicles' boundaries so that, to qualify as a vehicle, the target needs to have a sufficient support boundary on both sides of the vehicle.

When the configuration detection module is used as a pre-processor, all points in $\mathcal{L}$ at which templates match sets of edge pixels sufficiently closely are declared the centers of candidate targets.

Candidate points need to be further examined to verify that vehicles centered at these locations are geometrically consistent. The selection of competing points is carried out using textural information. During this stage, three textural features are computed for each candidate: (1) the mean value; (2) the standard deviation; and (3) the maximum of the intensity distribution of the pixels lying within the target-sized template. When two hypothesized vehicles are geometrically inconsistent, e.g., overlapping, the one having the higher quality measure survives. The quality measure is a weighted sum of the quality of the contour match and the similarity of the statistics to the reference statistics. The reference statistics are computed based on those of highly matched targets. The similarity of statistics is calculated by inverting a quadratic distance measure. At this step, most of the false alarms are rejected due to either a large quadratic distance or bad contour quality.

3.3 Verifier

The verifying stage primarily involves shadow verification. All recognition techniques suffer from the ambiguity existing between similar 2-D shapes arising from the contours of different 3-D objects. One way to avoid this problem is to use shadow information. Since shadows only accompany 3-D objects, false alarms can be minimized by checking for the existence of a shadow within a predicted region. A shadow has a high contrast and darker intensity than the background, and its size can be predicted using the vehicle orientation, viewing direction, and illumination direction. In our implementation, we use an enhancement process and a region growing process that segments an image into regions based on the intensity differences between connected components (Rodriguez, 1994). The resulting region segments are used to verify the existence of vehicle shadows. Detection is declared if there exist segments that comprise a region satisfying the following constraints: *Size constraint*: The size is computed by projecting the object model onto the ground plane. The portion of the shadow visible from the viewing direction is then computed. *Intensity constraint*: The shadow region is both homogeneous and dark. In other words, both the deviation and the mean value of the intensity distribution are small. *Position constraint*: The shadow position is consistent with the illumination direction. *Shape constraint:* The shadow can be on either one or two consecutive sides of the vehicle.

3.4 Detection on Long Roads

IU algorithms often fail to consider actual operational conditions where very large images are used. These images call for specific computational and sometimes algorithmic solutions. We have addressed this problem in the context of vehicle detection on long roads. To this end, site model information in the form of the road model is fed back to the low-level processing stages. More specifically, the Canny edge detection module is modified to perform filtering, gradient computation, non-maxima suppression, and hysteresis thresholding only on masked regions. Further improvements in speed using block processing of image chips extracted from the roads are being considered. Lastly, template specifications for a long road will take into account the fact that vehicles are often isolated, while some edges in parking lots are occluded.

3.5 Experiments

Detection results on Ft. Hood and TEC parking lots are given in Figures 1 and 2. The vehicle detector is used to infer historical trends, detect changes, or to implement negation mechanisms by which an IA can query the earliest date on which a significant event occurred (for example, the number of vehicles exceeding a given threshold). Change detection is illustrated with Ft. Hood images in Figure 1. Figure 1(a) shows an old image of a parking lot. Figure 1(b) shows the vehicles detected in Figure 1(a). Figure 1(c) shows an image of the same parking lot taken at a later date in the new Ft. Hood site. Figure 1(d) shows the vehicles detected in Figure 1(c). Similar detection results are shown for the Denver site in Figure 2. Finally, results of the application of the long road monitor to the old and new Ft. Hood images are shown in Figure 3.

4. DETECTION OF VEHICLE GROUPINGS

This module can be used as a tool for globally assessing the degree to which a parking lot is full, or determining the presence of platoons or convoys along roads. It is primarily intended as an attentional mechanism. This tool is based on the detection of special vehicle configurations—more specifically, those configurations that exhibit a periodic behavior. This module is of particular interest since it enables one to look globally for activity in an entire area. It also may serve as a pre-processing step to the local vehicle detection scheme.

Partial or complete prior information in the form of site and object models are used to infer information about the object geometry, the object grouping configuration geometry, or the hypothesized locations or ROIs that may include such configurations. Generic ROIs are either 1-D, and modeled as ribbons (e.g. road, railroad, landing strip, river, etc.), or 2-D, and modeled as regions (e.g., parking lot, open area, etc.). The detection of regular and periodic structures are carried out either by spectral analysis of the brightness intensity directional derivative in the 1-D case, or analysis of the gradient magnitude in the 2-D case. In the case of ribbons, the directional derivative is computed. Evidence of a specific periodic structure is found by searching for a peak within a frequency window, referred to as a *compliance window* derived from known object model constraints. The 2-D case is similar but yields *compliance regions* that entail the use of various search strategies depending on the type of prior information available about the configuration geometry and object geometry. The observation space for the detection rule includes the peak absolute spectrum magnitude and its value relative to its median computed within a specified frequency window. Acceptance regions are derived by using a decision-theoretic approach by exploiting a set of training images. These ideas are expanded and formalized in this section.

4.1 Characterization

Consider first the 1-D case where ROIs are represented by *ribbons* in the image. Let $\mathbf{r}$ denote the image position of a point. Assume that a ribbon's skeleton describes a smooth curve $\mathcal{C}$, given by $\mathbf{r}_{\mathcal{C}}(s)$, the image point position as a function of the curvilinear abscissa s. Let $\mathbf{N}(s)$ and $\mathbf{T}(s)$ be the normal and tangent vectors to $\mathcal{C}$ at $\mathbf{r}_{\mathcal{C}}(s)$. The ribbon R characterized by width $W(s)$ at $\mathbf{r}_{\mathcal{C}}(s)$ is defined by the set of points satisfying

$$R = \{\mathbf{r} : \mathbf{r} \in L(s)\}$$

with L(s) the set defined as

$$\{\mathbf{r} : \mathbf{r} = \lambda \mathbf{N}(s) + \mathbf{r}_{\mathcal{C}}(s), \lambda \in (-\frac{W(s)}{2}, \frac{W(s)}{2})\}.$$

Construct $k(s)$, the directional derivative $D_{\mathbf{T}}I(\mathbf{r})$ of the brightness intensity $I(\mathbf{r})$ averaged along the width of the ribbon:

$$k(s) = \int_{\mathbf{r}\in L(s)} d\mathbf{r}\, D_{\mathbf{T}}I(\mathbf{r})/W(s) \cdot$$

Periodically organized objects are detected from spectral analysis of directional edges $k(s)$. Let $K(f) = F[k|f]$ denote the Fourier transform (FT) of $k(s)$. Consider the occurrence of periodically situated objects, such as a road convoy, vehicles parked in a parking lot, or railroad cars on tracks. These are represented by a replicated function $c(.)$ – derived from the directional derivative profile of the object – within a window of length p, as

$$k(s) = \mathrm{rect}(\frac{s}{p})([c(.) * \sum_{n=-\infty}^{+\infty} \delta(. - n/f^*)](s)), \tag{1}$$

with $\mathrm{rect}(x) = 1$ if $|x| < 1/2$ and 0 if $|x| \geq 1/2$, for which $K(f)$ equals

$$K(f) = p \ \mathrm{sinc}(. \times p) * (C(.)f^* \sum_{n=-\infty}^{+\infty} \delta(. - nf^*))\,(f)\cdot$$

Periodic edge structure is characterized by detecting peaks and their corresponding harmonics from the amplitude spectrum $|K(f)|$. The problem is to recognize this spectrum when an additive clutter noise component is present, i.e., for $k'(s) = k(s) + n(s)$. The strategy used reliably to measure this occurrence is described in Section 4.3.

The 2-D case allows for many possible object configuration scenarios (as shown in Figures 4 and 5) and is relevant for analyzing situations involving regular object configurations in open areas. Let $k(\mathbf{r}) = ||\nabla I(\mathbf{r})||$, the gradient magnitude at $\mathbf{r}$, and consider the 2-D-FT,

$$K(\mathbf{f}) = F[k|\mathbf{f}] = \int\int_{-\infty}^{+\infty} d\mathbf{r}\ k(\mathbf{r})\ \exp -j2\pi\mathbf{r}^T\mathbf{f},$$

with $\mathbf{f} = [f_0, f_1]^T$. Consider a grouping of objects with arbitrary replications and orientation (see Figure 6); their gradient magnitude is described by

$$k(\mathbf{r}) = \sum_{\mathbf{n}} c(\mathbf{R}_{\omega_n}\mathbf{r} - \mathbf{v}_n), \tag{2}$$

with $\mathbf{v}_n$ and $\mathbf{R}_{\omega_n}$ respectively denoting an arbitrary 2-D translation vector and rotation matrix. Since

$$F[x(\mathbf{A}\mathbf{r})|\mathbf{f}] = |\det(A)|^{-1}\,F[x(\mathbf{r})|\mathbf{A}^{-\mathbf{T}}\mathbf{f}]$$

then with $k(\mathbf{r})$ as in Equation (2) we have

$$K(\mathbf{f}) = \sum_{\mathbf{n}} \exp\left(-j2\pi\mathbf{v}_n^T\mathbf{R}_{\omega_n}\mathbf{f}\right) F[c(.)|\mathbf{R}_{\omega_n}\mathbf{f}]\cdot$$

If the objects satisfy periodic configurations, i.e., the directions and intervals of replication are given by the vectors $\mathbf{v}_0$ and $\mathbf{v}_1$ (see Figure 6), with

$$\mathbf{v}_n = i\mathbf{v}_0 + j\mathbf{v}_1,$$

and the objects have identical orientations, i.e., $\mathbf{R}_{\omega_n} = \mathbf{R}_\omega$, within a support area A, then $k(\mathbf{r})$ becomes

$$k(\mathbf{r}) = \mathcal{I}_A(\mathbf{r})([c(\mathbf{R}_\omega .) * \mathrm{comb2}_{(\mathbf{v}_0,\mathbf{v}_1)}(.)]\,(\mathbf{r})),$$

with $\mathcal{I}_A(\mathbf{r})$ denoting the indicator function over A, $\mathrm{comb2}_{(\mathbf{v}_0,\mathbf{v}_1)}(\mathbf{r}) = \sum_{\mathbf{n}} \delta(\mathbf{r} + \mathbf{V}\mathbf{n})$, and $\mathbf{n} = [i,j]^T$. Its resulting spectral representation is (Dudgeon and Mersereau, 1984)

$$K(\mathbf{f}) = [I_A(.) * (C(\mathbf{R}_\omega .)\,\mathrm{Comb2}_{(\mathbf{v}_0,\mathbf{v}_1)}(.))]\,(\mathbf{f}),$$

where $I_A(.) = F[\mathcal{I}_A|\mathbf{f}]$, and $\mathrm{Comb2}_{(\mathbf{v}_0,\mathbf{v}_1)}(\mathbf{f})$ satisfies

$$\mathrm{Comb2}_{(\mathbf{v}_0,\mathbf{v}_1)}(\mathbf{f}) = \frac{1}{\det(\mathbf{V})} \sum_{\mathbf{n}} \delta(\mathbf{f} - (\mathbf{V}^{-1})^T\mathbf{n}).$$

where V is the 2-D periodicity matrix with columns composed by the $\mathbf{v}_i$'s, $i = 0, 1$, with $\det(\mathbf{v}_0, \mathbf{v}_1) = \det(\mathbf{V}) \neq 0$. This situation describes, by simple selection of the shape of the support region A, a wide array of configurations of special interest such as circular, semicircular, wedge, pyramid, block or linear formations (see Figure 7). As in the 1-D case, a periodic configuration can be detected by searching for its corresponding 2-D base and harmonic peaks whose positions are given by the frequency domain periodicity matrix $U = V^{-T}$. The rule for reliably detecting such peaks in the presence of noise is described in Section 4.3.

4.2 Spectral Searches on Compliance Windows

Constraints on 3-D geometrical objects and configuration models yield spectral compliance windows where dominant components characteristic of certain types of regular configurations can be searched. Let f^* denote the base component arising as a result of a regular object formation on a ribbon. Knowledge of the object geometry and context enables us to derive bounds on the positions of the peaks: Knowledge of the average object length l taken on the image region covered by the ribbon is derived from the known object model and the camera orientation parameters. Context also allows for knowledge of the minimum $(1+\epsilon)l$ as well as maximum $(1+\mu)l$ replication distance of this object. f^* must lie within a *spectral compliance window* W_c, defined as $W_c = [\frac{1}{(1+\mu)l}, \frac{1}{(1+\epsilon)l}]$. A regular structure is hypothesized if a peak is found at frequency f^*, also referred to here as the base or dominant spectral component:

$$f^* = \mathrm{argmax}(|K(f)|), \quad \text{where } f \in W_c.$$

When the dominant component is found in the spectral compliance window, corroborating evidence of harmonic components at $nf^*, n = 2, 3, \ldots$ is verified. The decision rule is built on a 2-D observation space composed of the logarithm of the spectrum magnitude and the logarithm of the spectrum magnitude normalized by the median (see Section 4.3). Similar compliance windows can be derived in the 2-D case (see Figure 8). Consider the periodicity matrix decomposition

$$\mathbf{V} = \hat{\mathbf{V}}\,\mathbf{diag}(\tau_0, \tau_1), \;\; \hat{\mathbf{V}} = [\hat{\mathbf{v}}_0|\hat{\mathbf{v}}_1]\,,$$

with $|\hat{\mathbf{v}}_0| = |\hat{\mathbf{v}}_1| = 1$. The spectral periodicity matrix $\mathbf{U}$ equals $\mathbf{U} = \mathbf{V}^{-T}$, so that $\mathbf{U} = \hat{\mathbf{V}}^{-T}\mathbf{diag}(\tau_0^{-1}, \tau_1^{-1})$. Note also that $\mathbf{u}_i.\mathbf{v}_j = \delta_{i,j}$, $i, j = 0, 1$ since $\mathbf{U}^T\mathbf{V} = (\mathbf{V}^{-T})^T\mathbf{V} = \mathbf{I}$. The spectral periodicity vectors $\mathbf{u}_i$ are orthogonal to their *alternate* spatial counterparts, since $\mathbf{u}_i.\mathbf{v}_j = 0$ if $i \neq j$. Using $\mathbf{u}_i.\mathbf{v}_i = 1$ and $\mathbf{v}_i = \tau_i\hat{\mathbf{v}}_i$, we get $|\mathbf{u}_i| = |\cos(\mathbf{u}_i, \mathbf{v}_i)|^{-1}\tau_i^{-1} = |\sin(\mathbf{v}_j, \mathbf{v}_i)|^{-1}\tau_i^{-1}$, with $j = \bar{i}$, so that if we let $\mathbf{u}_i = \hat{\mathbf{u}}_i\xi_i$, then ξ_i equals

$$\xi_i = \tau_i^{-1}\,|\sin(\mathbf{v}_0, \mathbf{v}_1)|^{-1}, \tag{3}$$

where it was assumed in the previous section that $det(\mathbf{v}_0, \mathbf{v}_1) \neq 0$. The case where $det(\mathbf{v}_0, \mathbf{v}_1) = 0$ consists of *linear* configurations in a 2-D domain, with a unique direction of replication, and one may write the gradient magnitude map as

$$k(\mathbf{r}) = \sum_n c(.) * \delta(\mathbf{r} - n\mathbf{v}),$$

for which

$$K(\mathbf{f}) = C(.)\sum_l \delta(\mathbf{v}^T\mathbf{f} - l),$$

which gives rise to impulsive lines orthogonal to the direction of $\mathbf{v}$. In sum, the spectral signature offers clues to the spatial geometry (orientation and dimension) of the detected configuration, and can be used to determine the spatial orientations of the object configurations. Alternatively, if site model information is available, compliance windows can be derived to guide the detection process. Using geometrical considerations similar to the ones presented in the 1-D case, the following compliance (longitudinal and lateral) constraints may be inferred:

$$\begin{gathered}\frac{1}{(1+\mu_0)l_0} < \mathbf{f}^{*T}\hat{\mathbf{u}}_0\,|\sin(\mathbf{v}_0, \mathbf{v}_1)| < \frac{1}{(1+\epsilon_0)l_0}, \\ \frac{1}{(1+\mu_1)l_1} < \mathbf{f}^{*T}\hat{\mathbf{u}}_1\,|\sin(\mathbf{v}_0, \mathbf{v}_1)| < \frac{1}{(1+\epsilon_1)l_1}\end{gathered} \tag{4}$$

where $l_0, l_1, \mu_0, \mu_1, \epsilon_0, \epsilon_1$ are derived from the task at hand. How much is known about V depends on the context. The search strategy for detecting a compliant impulsive component makes use of various schemes depending on the available model information.

- If only the orientation direction of the object configuration is available, then the constraints in Equation (4) translate into radial search spaces on a half-plane along directions $\hat{\mathbf{u}}_0$ and $\hat{\mathbf{u}}_1$, or alternatively into a search sector if an uncertainty on these directions is specified.

- If only the object dimensions are known, the search is conducted along annular rings, generated by a full 2π revolution of the sets in Equation (4).

- Finally, if both dimensions and directions can be inferred from the site model or from the context, the search space reduces to annular sectors.

4.3 Detection

Various factors can impair the detection process: The clutter noise component present across the spectral compliance window, poor imaging and illumination that lead to weak gradient magnitude, and inaccuracies in site and region boundary models. The factors interfering with the existence of a uniform critical region across experimental conditions are variations in viewing conditions and illumination.

We designed a detection rule that tests the dominant component at the base and corresponding harmonic frequencies. This rule takes into account the spectrum magnitude $|K(f^*)|$ associated with the maximum peak at f^* within a compliance window. To keep critical/detection regions constant across viewing directions, the spectral estimates need to be normalized by the aspect ratio. The resulting normalized spectrum magnitude associated with the peak is denoted by K_a. The second element in our decision rule is the ratio of the spectrum magnitude over the median of this magnitude computed over the compliance window: $K_r(f^*) = |K(f)|/K_{\mathrm{med}}$, with K_{med} denoting the "median" magnitude inside the spectral compliance window. The detection acceptance and rejection regions are designed over the 2-D observation vector $\mathbf{Y}$ whose components are the logarithms of these two measures, i.e., $\mathbf{Y} = (\ \ln(K_a), \ln(K_r)) = (L_a, L_r)$. The harmonics nf^* or $\mathbf{n}^T\mathbf{f}^*$ are also checked for the presence of an impulsive component. A decision-theoretic technique used to determine the optimal critical region on a control set of images is described in detail in Section 5.6.

4.4 Application

After registration, the site model information is used in the newly acquired image. ROIs are delineated and image object dimensions are derived. Figure 9(a) shows an initial image, and Figure 9(b) shows the new image precisely registered to the site model. From the camera model, 3-D models of ROIs can be transformed and registered to the image plane. In Figure 9(c), roads, parking, and staging areas are delineated. The attentional mechanism and local object detection schemes are applied to the monitored areas highlighted in Figure 9(d).

As previously explained, compliance window parameters are derived from context information. For example, if the task is convoy detection on roads, the average image vehicle length l taken on the image region covered by the road is first derived from the known vehicle model and the camera orientation parameters. For safety reasons, vehicles follow each other at distances greater than ϵl, for some $\epsilon > 0$; an upper bound for f^* is then given by $1/((1+\epsilon)l)$. Similarly, vehicles following each other at too great distances are not considered to belong to convoys. The maximum distance between successive vehicles can then be expressed in terms of vehicle lengths l as $(1+\mu)l$ with $\mu > \epsilon$. This yields a lower bound $1/((1+\mu)l)$ for the dominant spectral component and for the minimum length. Similarly, when dealing with monitoring 2-D regions such as parking lots, partial or complete information can be derived from the site model about the longitudinal and lateral bounds $l_0, l_1, \mu_0, \epsilon_0, \mu_1, \epsilon_1$ as well as the orientation yielding the periodicity matrix. These correspond to vehicle dimensions as well as spacing between neighboring vehicles in parking lots. The various degrees to which this information is known yields the specific search strategies described in the previous section. In the case of train detection, these constraints are even more precise, and can be readily derived from the types of trains being monitored (freight, passenger, etc.).

The above detection scheme can be applied directly to convoy detection or the detection of full parking areas. Alternatively, the above mechanism can serve as a pre-processing attentional mechanism prior to the application of more computationally intensive detection and counting algorithms. We use it to cue a local vehicle detector. To this end, we infer the spatial positions of edges giving rise to the periodic structure determined through frequency domain analysis, referred to here as periodic loci, and thereafter, apply local object detection schemes to corroborate the presence and refine the count of objects at these positions. Let $\mathcal{L}$ denote the set of periodic loci. The local vehicle detector can then be applied to the reduced spatial window centered at these locations in $\mathcal{L}$.

4.5 Experiments

Our first experimental results were obtained using a set of 40 images from the RADIUS-MB2 series. While this is a modest test sample, these images are interesting since they offer an array of varying viewing and illumination conditions as well as activity scenarios. Figure 9(a) shows an original image (M17). In Figure 9(b), the site model is superimposed on the image. After registration of M17 (Chellappa *et al.*, 1994), a selected number of delineated regions of interest (roads, parking areas, and staging areas) are derived from the camera external orientation parameters and the site model, as shown in Figure 9(c). Figure 9(d) shows the parking and road areas monitored in our experiment. The needed apparent image dimensions of monitored objects, as well as road and parking area boundaries and directions, are derived from the known site model and resected camera parameters for the image under analysis.

In the 1-D detection case, directional derivatives are calculated on ROIs made up of ribbons. For roads, this direction is the direction of the tangent to the road; for 1-D parking areas, this direction is the dominant orientation of the parking area.

For images M10, M18, M20 and M33, the resulting detected active roads are highlighted in Figures 10(a), (b), (c), and (d). Detected active parking areas are highlighted in Figures 11(a), (b), (c), and (d) for images M10, M1, M21 and M11, respectively. The resulting set of periodic loci $\mathcal{L}$ is shown in Figure 12(a). After application of the vehicle detector, the resulting detected vehicles are shown in Figure 12(b). Visual inspection shows that the results are consistent with the actual images.

Detection and false alarm rates resulting from running the global activity detector on all or a subset of the 40 model board MB2 images are reported here for hand-tuned rules (results for learned decision rules are reported in Section 5.6). The detection of convoys is tested on some of these images using a tuned detection rule, yielding a compound non-detection probability of 0.10 and a false alarm probability of 0.16. The detection of active (full) parking lots on all the MB2 images yields a compound non-detection probability of 0.07, and a false alarm probability of 0.17 for hand-tuned critical region boundaries. The breakdown of the non-detections and false alarms is provided for the MB2 images in Tables 1 and 2 for convoy detection and parking lot occupancy analysis.

The previous result shows that the detection rule is not too sensitive to the varying illumination and viewing conditions present in the dataset. We wish to better characterize the sensitivity of the detection performance to misspecification of the object dimensions. To this end, the 3-D dimensions (width and length) of the vehicle were varied, and the compound detection and false alarms were computed on the set of test images. The results are reported in Section 5.6.

For specific applications such as those described here, this method compares favorably with the direct application of symbolic recognition schemes. By using a global operator we can readily identify specific ROIs, while the complexity of this method is minimal compared to symbolic recognition. Delineation from known site models, computation of the gradient, and computation of the spectrum are of modest computational complexity. These combined operations take approximately 10 seconds for an image of size 1320 by 1035 on a SUN SparcStation^{TM}.

The same process is carried out in 2-D image regions. The global analysis is applied to an image taken from the Ft. Hood image series shown in Figure 13. An enlarged area is shown in Figure 14 for which the gradient magnitude is computed. We show the corresponding edge map obtained by thresholding in Figure 15. This image exhibits two types of periodicities; one corresponds to the row periodicity, the second to the car periodicity. The "principal" direction is parallel to the spatial row periodicity direction. The vehicles detected by the local detector are shown in Figure 16. A contour plot of the spectrum with additional level sets is displayed in Figure 17. The base frequency corresponding to the vehicle periodicity is searched for along a direction orthogonal to the principal direction, since the image is a nadir view. This leads to a directional tolerance area whose boundaries are indicated in Figure 18. On the other hand, geometric constraints yield a compliance window in the form of an annular ring. The base frequency corresponding to the vehicle periodicity lies at the maximum of the spectrum in a compliance window equal to the intersection of the annular ring and the angular sector, as shown in Figure 18. For this experiment $\epsilon = 1.2$ and $\mu = 5$ were chosen.

Inferring the orientations is interesting both for registration and site model construction purposes. If these orientations are included in the site model, this information allows the refinement of camera exterior orientation parameters. If this information is not already contained in the site model, it is relevant since it may be used as a site model construction tool. The principal and secondary directions, inferred from searching for the maximum and second largest peaks in an annular search area (itself determined from vehicle dimension constraints), leads to spatial orientations that are correct to within 4 degrees when compared to the orientations measured from the actual images.

5. SENSITIVITY ANALYSIS AND LEARNING

Part of our work has focused on providing sensitivity analysis and learning strategies for the movable object detection modules. In particular, we have considered the sensitivity of the tuning parameters that characterize the vehicle detection and counting algorithms that we have integrated into the RADIUS testbed.

5.1 Procedural Approach

Robustness as well as the ability to work in an unsupervised mode are two desirable features for the algorithms developed in the cont
text of the RADIUS project. A more modest yet still important goal that we have pursued involves the problem of automatic parameter tuning. Manual tuning is often neither a desirable nor a practical option because of the large amount and diversity of imagery typically used in exploitation. The need for limiting the number of tuning parameters is further justified by the fact that IAs should not be required to deal with low-level parameter tuning.

While the problem of automatic parameter selection will be partially addressed by the use of a rule-based system such as the HUB (Strat, 1995), we are addressing the problem of limiting or completely eliminating the number of tuning parameters using the following strategy. Preliminary sensitivity analysis is used to identify the parameters to which the algorithm is the most sensitive. A compound measure of sensitivity is chosen as the expected risk function and is computed empirically over a set of test sites. The parameters to which the algorithm is the least sensitive are "frozen" to their optimal values (off-line parameter optimization). We provide on-line training tools for the remaining parameters on which performance depends the most.

The approaches used for parameter training and optimization are set in a classic hypothesis testing framework using Bayesian or Neyman-Pearson strategies.

A related issue is the assessment of the sensitivity of the algorithms to misspecification of model parameters. This problem is important for model-based exploitation and some initial results are reported here.

5.2 Sensitivity to Tuning Parameters

The sensitivity of the vehicle detection and counting module with respect to the edge detection step has been extensively studied. Specifically, we characterize its sensitivity to the Canny parameters: mask size and low and high thresholds. We show in Tables 3, 4, and 5 the detection and false positive rates for varying values of the Canny mask size as well as the low and high thresholds, for detection and counting on the parking lots in Ft. Hood and Denver TEC2. The values reported in Tables 4 and 5 correspond to the false positive rate[1] and detection probability optimized over the mask size. In turn, the

[1] For vehicle detection and counting, a false positive rate is computed instead of a probability of false alarm; this rate is the ratio of the number of wrongly detected vehicles to the total number of detected vehicles. This measure of false alarm differs from the probability of false alarm. An alternate false positive rate often used is the rate per area processed. For detection of convoys, the usual probability of false alarm is used instead.

low and high thresholds are set to their optimal values (in the sense described in Section 5.4) to produce the false positive rate and detection probability as functions of mask size in Table 3.
Next we characterize sensitivity to the overlap threshold. Table 6 shows the false positive rate and detection probability. From these experiments we compute the empirical expected risk $E\{R(d)\}$ including both false alarm and non-detection rates. These are presented in Sections 5.4 and 5.6. From this empirical risk it can be inferred that the variation in detection performance is smaller for the Canny parameters when they vary within their operational limits, while the performance varies significantly with the overlap threshold. For the former parameters we need to derive optimal estimates; for the latter parameter, we need to provide a training tool.

5.3 Detection and Training

On-line parameter training and off-line parameter optimization are set in a hypothesis testing framework. Let H_0 and H_1 correspond to the two hypotheses (absent/present),

$$H_i : P(\mathbf{Y}|i) = P_i(\mathbf{Y}).$$

The acceptance and rejection regions are designed over some observation space $\mathcal{O}$. The observation vector $\mathbf{Y}$ in the case of the vehicle detector is simply the overlap threshold. In the case of the configuration detector, it is composed of two spectral measures described in Section 5.6.

A decision rule d is simply given by $d(\mathbf{Y}) = \mathcal{I}_{\mathcal{R}}(\mathbf{Y})$, with $\mathcal{I}_{\mathcal{R}}$ the indicator function on the acceptance region $\mathcal{R}$. The acceptance region admits a parametric form

$$\mathcal{R} = \mathcal{R}_{\mathcal{V}} = \{\mathbf{Y}\text{such that } \ b(\mathbf{Y};\mathcal{V}) \geq 0\},$$

where the form of the critical/acceptance regions can be inferred from the distributions satisfied by the observation vector under either hypothesis. The design of the detection rule is set up according to the following two strategies:

- (A) Bayesian strategy: Find $\mathcal{V}^* = \text{argmin}(E\{R(d)\})$. This strategy consists of minimizing the expected risk (Poor, 1988) where the cost factors C_{fa} and C_{nd} are chosen to balance the cost associated with a false alarm and a non-detection:

  $$E\{R(d)\} = C_{nd}P_0(\mathcal{R}_{\mathcal{V}}).\pi_0 + C_{fa}P_1(\mathcal{R}_{\mathcal{V}}{}^c)\pi_1,$$

  where $P_0(R)$ and $P_1(\mathcal{R}_{\mathcal{V}}{}^c)$ are the false alarm and non-detection probabilities, and π_i are the priors.

- (B) Neyman-Pearson strategy: Find $\mathcal{V}^* = \text{argmax}(P_1(R))$ subject to $P_0(R) < \alpha$.

These above strategies are used as follows:

- **Canny Parameter Optimization**: For determining optimal Canny parameters for vehicle detection we use (A) where $\mathcal{V}$ includes the mask size and threshold parameters and the empirical expected risk is computed over a training set.

- **Overlap Threshold Training**: For the overlap threshold in vehicle detection we use (B) where $\mathcal{V}$ is the threshold and $P_1(\mathcal{R}_{\mathcal{V}}{}^c)$ and $P_0(\mathcal{R}_{\mathcal{V}})$ are computed empirically from a user-provided reference patch.

- **Convoy Detection Training**: For determining the acceptance region for detection we use (A) where $\mathcal{V}$ parameterizes the region boundary in the observation space and $E\{.\}$ is empirically computed over a training set.

5.4 Parameter Training and Optimization for the Detection and Counting Module

For off-line Canny parameter optimization we use a Bayesian strategy (A) with the expected risk computed over the training set. We choose $C_{nd} = 0.50$ and $C_{fa} = 0.50$ and we assume equal priors, which is equivalent to a minimum probability of error scheme. The empirical value of $E\{R(d)\}$ computed over the range of Canny parameters is shown as a function of the Canny threshold and mask sizes in Figures 19 and 20. The Canny parameters are set off-line to the values minimizing the expected risk. These values were found to be 5 for the mask size and $(200, 400)$ for the minimum and maximum thresholds.

The expected risk computed for varying values of the overlap threshold shows that the performance depends more closely on this parameter (Figure 21). For on-line overlap threshold training, we use strategy (B). Empirical detection and false alarm probability are derived from predetermined reference patches. For each newly acquired image, the overlap threshold is automatically computed from the empirical probabilities derived from the reference patches. Denote by $P_1(\mathcal{R}_{\mathcal{V}}) = r(P_1(\mathcal{R}_{\mathcal{V}}{}^c))$ the empirical ROC curve derived by varying the overlap threshold $\mathcal{V}$; then we want to find $\mathcal{V}^*$ satisfying $\mathcal{V}^* = \text{argmax}(P_1(\mathcal{R}_{\mathcal{V}}))$ subject to

- $P_0(\mathcal{R}_{\mathcal{V}}) < \alpha$,

- $\frac{d(r(p))}{dp}_{\,p=P_1(\mathcal{R}_{\mathcal{V}}{}^c)} > \beta$.

The second condition ensures that the slope of the ROC curve for the given value of $\mathcal{V}$ is not too small, i.e., that an increase in false alarm probability can be traded for a significant increase in detection probability. Several scenarios are considered in the next section for the selection of reference patches.

5.5 Selection of Reference Patches for the Detection and Counting Algorithm

When a new image is acquired, the overlap threshold is automatically "calibrated" using reference patches as explained above. In certain situations, reference patches with their associated ground truth can be determined by inspection, and the ground truth can be represented as site features in the site model. This is possible when we can identify patches in the parking lot that are known to be always empty (such as passageways, parking lot exits, etc.), or always full (areas located near building entrances). If such patches cannot be reliably identified, an alternate procedure consists of using fixed structures such as buildings for calibrating the overlap threshold. Lastly, a more interactive method would consist of letting the IA inspect and validate the reference patches over a subset of the images to be processed before initiating a batch procedure.

5.6 Parameter Training for Convoy Detection

The convoy detection rule tests the dominant spectral component within a compliance window at the base and corresponding harmonic frequencies, and takes into account the normalized spectrum magnitude K_a associated with the maximum peak at f^* within a compliance window and the ratio of the spectrum magnitude to the median K_{med} of this magnitude computed over the compliance window. Specifically, the components of $\mathbf{Y}$, the 2-D observation vector, are the logarithms of these two measures, i.e., $\mathbf{Y} = (\ln(K_a), \ln(K_r)) = (L_a, L_r)$. Let H_0 and H_1 correspond to the two hypotheses, with H_0 the hypothesis that no peak is present; the decision rule is simply $d(L_a, L_r) = \mathcal{I}_{R_\mathcal{V}}(L_a, L_r)$. We use a Bayesian strategy for deriving the acceptance region from a training set of images. The acceptance region boundary is parameterized by vector $\mathcal{V}$ and chosen as

$$R = R_\mathcal{V} = \{(L_a, L_r), \text{such that}\, b(L_a, L_r; \mathcal{V}) \geq 0\}.$$

Assume that the joint conditional probability distributions on $\mathbf{Y} = (L_a, L_r)$ are Gaussian, i.e., $H_i : P(\mathbf{Y}|i) \sim N(\mathbf{m}_i, \Sigma_i), i = 0, 1$; then the log-likelihood ratio function is a quadratic function in $\mathbf{Y}$(Fukunaga, 1990), i.e. $(\mathbf{Y} - \mathbf{m}_1)^t \Sigma_1^{-1} (\mathbf{Y} - \mathbf{m}_1) - (\mathbf{Y} - \mathbf{m}_0)^t \Sigma_0^{-1} (\mathbf{Y} - \mathbf{m}_0)$. We assume dissimilar covariances[2] for which the boundary equation $b(.,.;.) = 0$ is a conic section. The acceptance region is determined by finding $\mathcal{V}^*$, which minimizes the expected value of the conditional risk computed over the training set, i.e., $\mathcal{V}^* = \text{argmin}(E\{R_\mathcal{V}(d)\})$. Ten of the forty MB2 images were chosen as a training set, and $b(.,.;.)$ was assumed to be an elliptic boundary. The parameters of this elliptic boundary were optimized on the set of control images[3]. The resulting boundary for $C_{nd} = 0.55$ and $C_{fa} = 0.45$ is shown in Figure 22. When 10 of the 40 images were chosen as a training set, the compound detection performance when compared to the hand-tuned performance yielded a significantly better false alarm probability, 0.11, and a comparable non-detection probability of 0.08 (see Table 7). These results can be contrasted to those reported in Table 2.

5.7 Sensitivity to Misspecification of Model Parameters

We characterize the sensitivity of parking lot occupancy detection to misspecification of the model parameters. The 3-D dimensions (width and length) of the vehicle are varied, and the detection and false alarm probabilities are computed on the set of test images. The resulting probabilities are displayed as functions of these dimensions in

[2] If the covariance matrices are equal, the boundary region is a hyperplane. If they are dissimilar, the boundary is elliptic, parabolic, or hyperbolic, depending on the covariance matrices.

[3] The expected value $E\{R(d)\}$, computed over the training set, is a noisy function of V, in part due to the modest size of the training set. V* is determined by using the Nelder-Mead Simplex algorithm (Nelder and Mead, 1965). This function is non-convex and, therefore, the simplex algorithm is not guaranteed to converge. Furthermore, the minimum is not unique.

Figure 23 for detection of active parking lots. In Figure 23, the upper surface represents the probability of detection as a function of the 3-D width W and length L. The lower surface represents the false alarm probability[4]. In this figure we see that the resulting performance is not too sensitive to a reasonable variation in size. Figure 24 shows the resulting expected risk computed for the above C_{nd} and C_{fa}. The situation is different for convoy detection, as seen from the simulation results in Figure 25. In this case, as the values deviate from their optimal specifications (the middle of the grid), performance degrades. As W and L increase, the false alarm probability decreases along with the detection probability. This highlights the importance of context as well as the adequate specification of model parameters for this particular application.

5.8 Qualitative Assessment of the Vehicle Detector Performance and the HUB

As a byproduct of this sensitivity analysis, we can infer the conditions for which the module works as expected, and relate these to the context in which the image was taken. This information can be encoded into the HUB as Prolog rules guiding the application of specific algorithms (Strat, 1995) according to image quality and context. Additional qualitative conditions for the proper application of the vehicle detector have been determined.

Based on multiple experiments carried out on the Ft. Hood and Denver images, we conclude that the current vehicle detector modules should be used with caution when any of the following conditions are present:

- Poor contrast
- Very cluttered environment or heterogeneous objects (such as some Ft. Hood storage areas)
- Vehicles of greatly varying orientations in a parking lot when only one orientation was specified
- Low image resolution such that the image dimension of a vehicle is ≤ 6 or 7 pixels
- Presence of water pools or oil marks.

These conditions could be coded into Prolog rules in an RCDE HUB subsystem that could control the appropriate application of detection algorithms.

[4] Situations where the width is greater than the length constitute a misspecification by $\pi/2$ of the actual vehicle orientation.

6. SYSTEM INTEGRATION OF THE VEHICLE DETECTION AND COUNTING MODULE

6.1 Exploitation Module

The RCDE platform implements basic photogrammetric procedures and low-level image processing and manipulation functions, and it provides a common development environment on which RADIUS-related algorithms can be developed and tested. It also provides a framework from which the RADIUS testbed system is derived. The vehicle detection and counting module presented here has been integrated into the RADIUS testbed and the remaining modules have been developed around the RCDE system.

A typical exploitation scenario making use of the RCDE basic functionalities is conducted as follows. A site model containing permanent objects is constructed. The prerequisites of the site model construction are the availability of several partially overlapping images of the same site. Multiple views of the site objects are necessary to disambiguate the object positions and dimensions. When images are simultaneously displayed, a shared 3-D local world coordinate system is defined for all images, and registration is carried out. As previously mentioned, in the usual photogrammetric applications, images are provided with a set of interior as well as exterior orientation parameters with associated covariance matrices. These are usually supplied in a TEC header format. These initial orientation parameters are refined by using our registration methods along with RCDE-supplied resection tools, after control point objects have been defined (these are points whose 3-D positions are known). Control points are used along with their identified image positions to refine camera orientation. Conjugate points are determined, either manually, or using one of our automated procedures. Single or multiple image resection is then accomplished.

After the image has been registered to a common local vertical coordinate system, site construction can be initiated. Three-dimensional generic objects are selected, created, and modified. Objects of primary importance in our case are roads, buildings, delineated training areas, and parking lots. Objects are manipulated and modified interactively by the operator, or defined through dialog windows. This site-generation procedure is carried out until a satisfactory site model has been constructed from all viewpoints.

When a new image is acquired, it is registered to the existing site model, and is thereupon exploited. Site models are not static. While the detection algorithms reported here exploit site model information to detect changes, they can be used, in turn, to verify and detect discrepancies, and therefore, to refine and update site model information. These algorithms are not indiscriminately applied to all incoming images and all areas. Instead, the IA specifies which areas should undergo special scrutiny.

One operational procedure of particular importance in driving the exploitation of images is determined by the concept of a quick look (QL) (Bailey *et al.*, 1994). In the QL mode, only small areas (usually corresponding to functional regions) are processed in batch over multiple images; historical comparisons are then carried out to determine trends and evolutions. Only a limited number of areas in the site are exploited; usually small changes in these areas are significant. Changes must, however, be detected in a timely fashion over large sets of images. Active QL areas are then prioritized for further and finer exploitation.

In the QL scenario, areas supporting the presence of convoys are brought to the attention of the IA. The IA may then trigger a specialized vehicle detector module, according to the type of ROI (road or parking). Exact vehicle counts are then reported to the analyst.

6.2 Integration Details

6.2.1 Assessing Operational Requirements for the Vehicle Detector

The vehicle detection and counting algorithm has been fully integrated into the RCDE platform. It was first implemented in C, and later re-implemented in C++. The RCDE platform is instead implemented in Lucid Lisp. Part of the vehicle detection process, consisting mostly of the interface, is implemented in Lisp. The Lisp/C interface of the RCDE enables inter-process communication between these two components. The integrated vehicle detector has been tested by several institutions, including GE and Lockheed-Martin Management and Data Systems. Its interface is currently being tailored to the specific needs of IAs.

Figure 26 shows an example of a dialog window for the integrated implementation of the vehicle detector module. After vehicles have been detected, the result is stored in the database for subsequent inspection.
The vehicle detector integrated into the RADIUS testbed currently includes the following features: The vehicle detector is accessible either stand-alone or under the Quick Look (QL) menus. It allows for dual specification of the vehicle dimensions: 2-D or 3-D. It enables a two-way interactive specification of the dimensions: Either graphically or by using sliding bars. The user can create a sample image vehicle, modify its dimensions, and then use it as a prototype. The user can modify its 3-D dimensions directly. The sample vehicle size is then modified accordingly. Alternatively, the user can modify the sample vehicle size, and the 3-D dimensions are modified accordingly. The detector accepts Fast-Block Interpolation Projection (FBIP) format and lazy images. It also is equipped with a ground truth verifier. It has been modified to accommodate some new function definitions in RCDE 3.0. It can handle batch processing in QL menus.

6.2.2 Integration into the QL Menus

Menus were modified so as to reflect recommendations made by NEL analysts. The detector was finally modified so as to allow operation under the newly designed LM QL menus, and to allow its application over large image databases in batch modes. Examples of QL menus are shown in Figure 27.

6.2.3 An Automatic Ground Truth Verifier

Performance statistics can be obtained by inspection. We have implemented an automatic ground-truthing tool to allow for the automatic assessment of the statistical performance of our vehicle detector. It is based on a tool originally developed by GE. This is a preliminary developer's tool allowing us to perform a fast evaluation of the results when a change in the vehicle detection scheme has been made. As a first approximation, the detection criterion declares a correct detection whenever the center of the detected vehicle is within the boundaries of a specified ground truth box. This rule will later be refined so as to dissociate the detection rate and the quality of the localization. The module allows for saving previous detection results and ground truth as site model objects. This is the preferred method for specifying the ground truth, since it can be reloaded and more easily exchanged.

7. CONCLUSIONS

While the movable object monitoring tasks addressed here were of limited scope, the problem has been addressed in all its facets. Our work has spanned concept design, implementation, transition to the RCDE and testbed platform, inclusion in QL operation, constant update and improvement of the algorithms, and finally, development of learning modules as well as sensitivity analysis. We believe this work, along with that of other groups involved in RADIUS, can serve as a preliminary model of how the transitioning of mature algorithms can be combined with the development of novel ideas.

The other accomplishment of this project has been to show that satisfactory reliability of algorithms can be obtained using context and model information. Future should be characterized by an increased emphasis on contextual information as well as the inclusion of temporal reasoning.

8. REFERENCES

ARPA. RADIUS testbed system operations concept, July 1993.

J.S. Bailey, M.D. Kelly, and J.D. Sargent. Quick-look: A new way to prioritize imagery for exploitation. In *Proc. ARPA Image Understanding Workshop*, pages 247–249, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

M. Bejanin, A. Huertas, G. Medioni, and R. Nevatia. Model validation for change detection. In *Proc. ARPA Image Understanding Workshop*, pages 287–294, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

J.F. Bronskill, J.S.A. Hepburn, and W.K. Au. Knowledge-based approach to the detection, tracking and classification of target formations in infrared image sequences. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 153–158, June 1989.

R. Chellappa, Q. Zheng, L.S. Davis, C.L. Lin, X Zhang, C. Rodriguez, A. Rosenfeld, and T. Moore. Site-model based monitoring of aerial images. In *Proc. ARPA Image Understanding Workshop*, pages 295–318, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

R.T. Collins, A.R. Hanson, and E.M. Riseman. Site model acquisition under the UMass RADIUS project. In *Proc. ARPA Image Understanding Workshop*, pages 351–358, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

D.E. Dudgeon and M.M. Mersereau. *Multidimensional Digital Signal Processing.* Prentice-Hall, 1984.

K. Fukunaga. *Statistical Pattern Recognition.* Academic Press, 1990.

A. Huertas and R. Nevatia. Detecting buildings in aerial images. *Computer Vision, Graphics and Image Processing*, 41:131–152, 1988.

J. Kawamura. Automatic recognition of changes in urban development from aerial photographs. *IEEE Trans. on Systems, Man and Cybernetics*, 1:230–240, 1971.

A.B. Khang and G. Robins. Optimal agorithms for extracting spatial regularity in images. *Pattern Recognition Letters*, 12:757–764, 1991.

D.M. McKeown and M. Roux. Feature matching for building extraction from multiple views. In *Proc. ARPA Image Understanding Workshop*, pages 331–349, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

D.M. McKeown. Toward automatic cartographic feature extraction. In *Mapping and Spatial Modeling for Navigation*, pages 149–180. Springer-Verlag, 1990.

J. Nelder and R.A. Mead. A simplex method for function minimization. *The Computer Journal*, 7, 1965.

V. Poor. *An Introduction to Signal Detection and Estimation.* Springer-Verlag, 1988.

K. Price. Symbolic matching of images and scene models. In *Proceedings of the Workshop on Computer Vision*, pages 105–112, 1982.

C. Rodriguez. An appearance-based approach to object recognition in aerial images. Master's thesis, Department of Computer Science, University of Maryland, College Park, MD, 1994.

SRI International and Martin Marietta. *RADIUS Common Development Environment:User's Manual*, July 1993.

T.M. Strat and W.D. Climenson. Site model content. In *Proc. ARPA Image*

Understanding Workshop, pages 277–285, Monterey, CA, 1994. Morgan Kaufmann, Publisher.

T.M. Strat. Integrating IU algorithms in the RADIUS HUB. Software documentation, 1995.

M.S. Ulstad. An algorithm for estimating small scale differences between two digital images. *Pattern Recognition*, 5:323–333, 1973.

V. Venkateswar and R. Chellappa. Extraction of straight lines in aerial images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16:1111–1116, 1992.

L. Xu. A method for recognizing configurations consisting of line sets and its application to discrimination of seismic face structures. In *Proc. International Conference on Pattern Recognition*, pages 610–612, 1988.

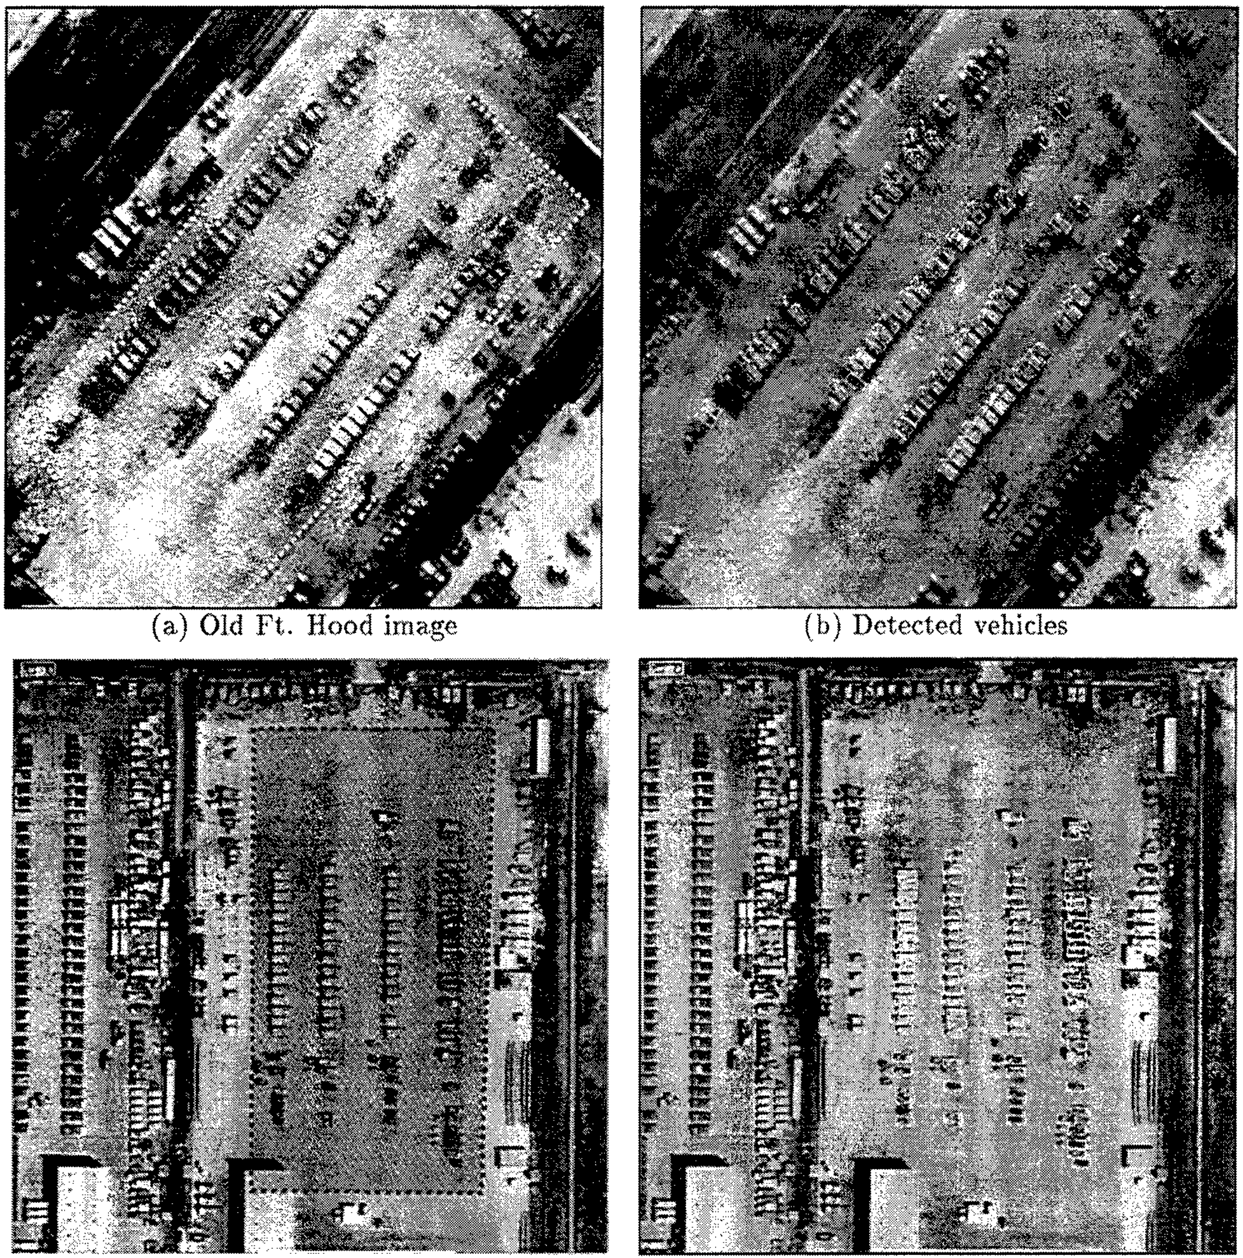

(a) Old Ft. Hood image

(b) Detected vehicles

(c) New Ft. Hood image

(d) Detected vehicles

Figure 1: Model-supported vehicle detection results on Ft. Hood images. (a) A parking lot is delineated in an old Ft. Hood image. (b) The detected vehicles. (c) The same parking lot on a new Ft. Hood image. (d) The newly detected vehicles support change detection and negation.

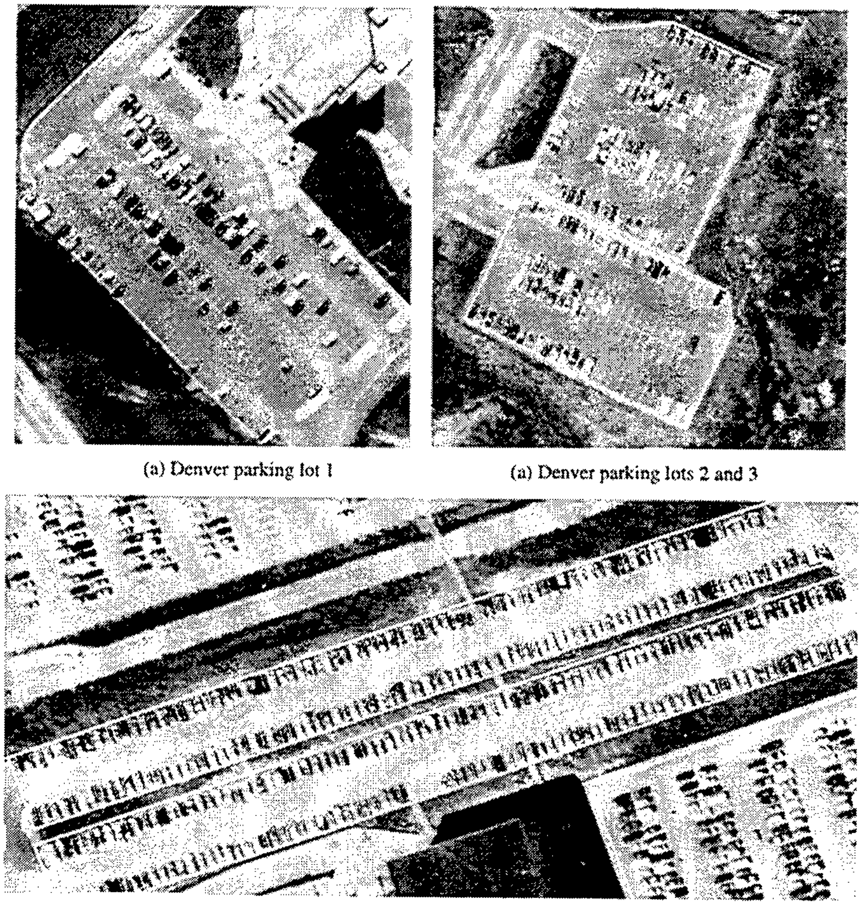

(a) Denver parking lot 1

(a) Denver parking lots 2 and 3

(c) Denver parking lots 5,6 and 7

Figure 2: Model-supported vehicle detection on Denver site parking lots 1,2,3,5,6, and 7.

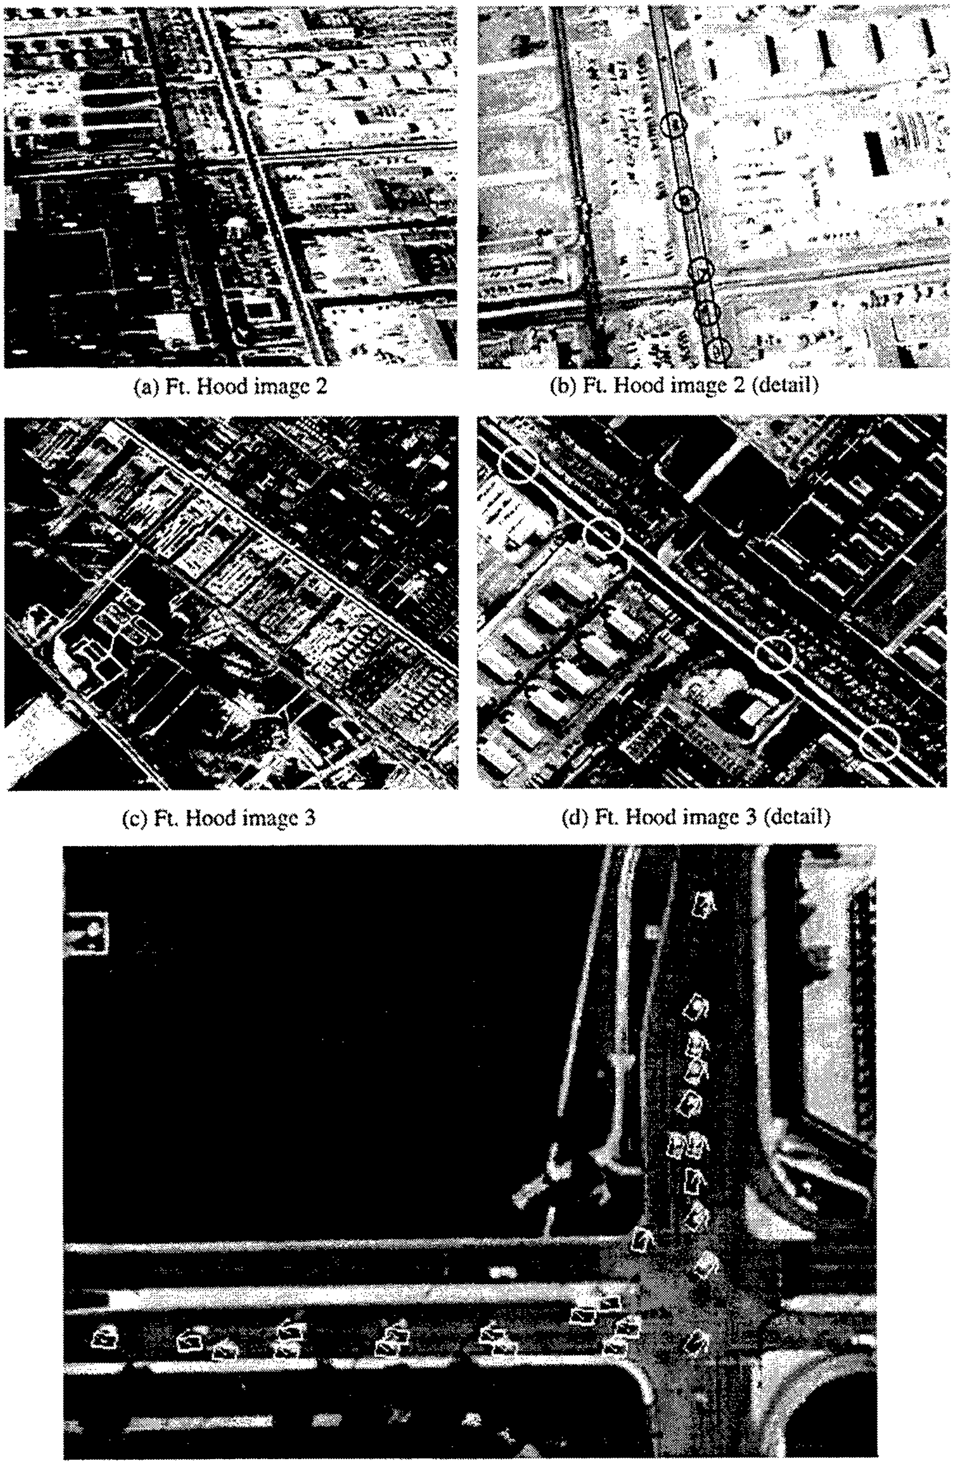

(a) Ft. Hood image 2

(b) Ft. Hood image 2 (detail)

(c) Ft. Hood image 3

(d) Ft. Hood image 3 (detail)

(e) New Ft. Hood image fhn711

Figure 3: Vehicle detection on roads for Ft. Hood images.

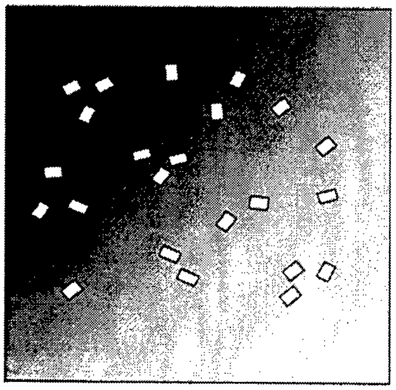

Figure 4: Arbitrary configuration.

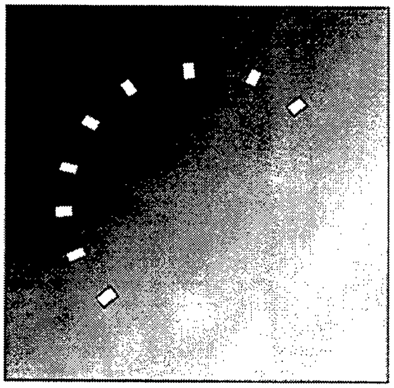

Figure 5: Circular configuration.

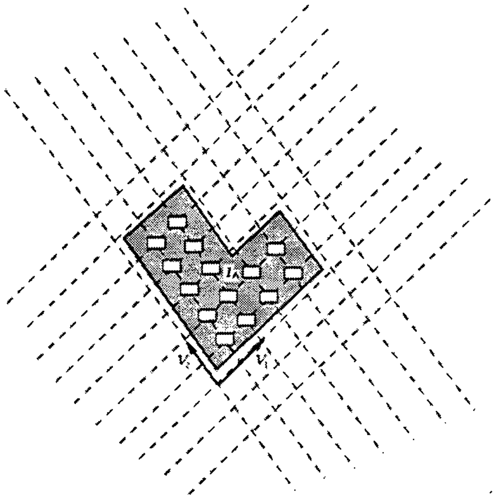

Figure 6: Canonical configuration representation.

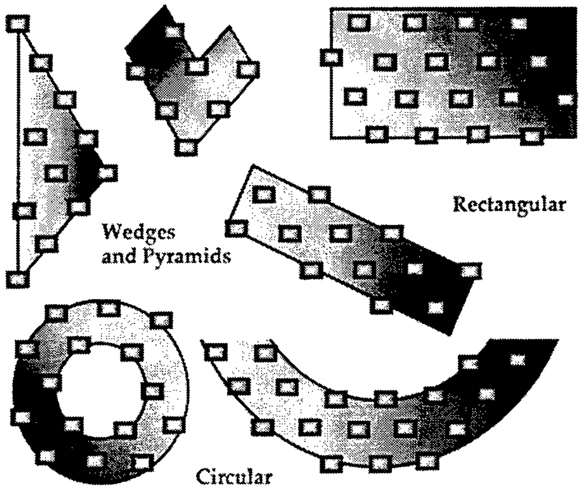

Figure 7: Periodic configurations.

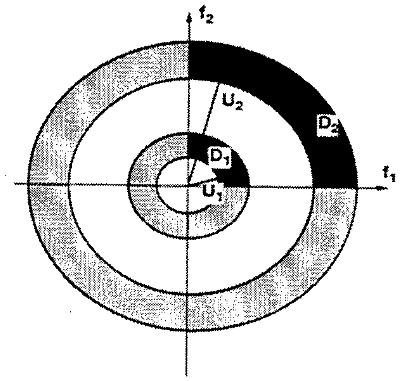

Figure 8: Compliance window.

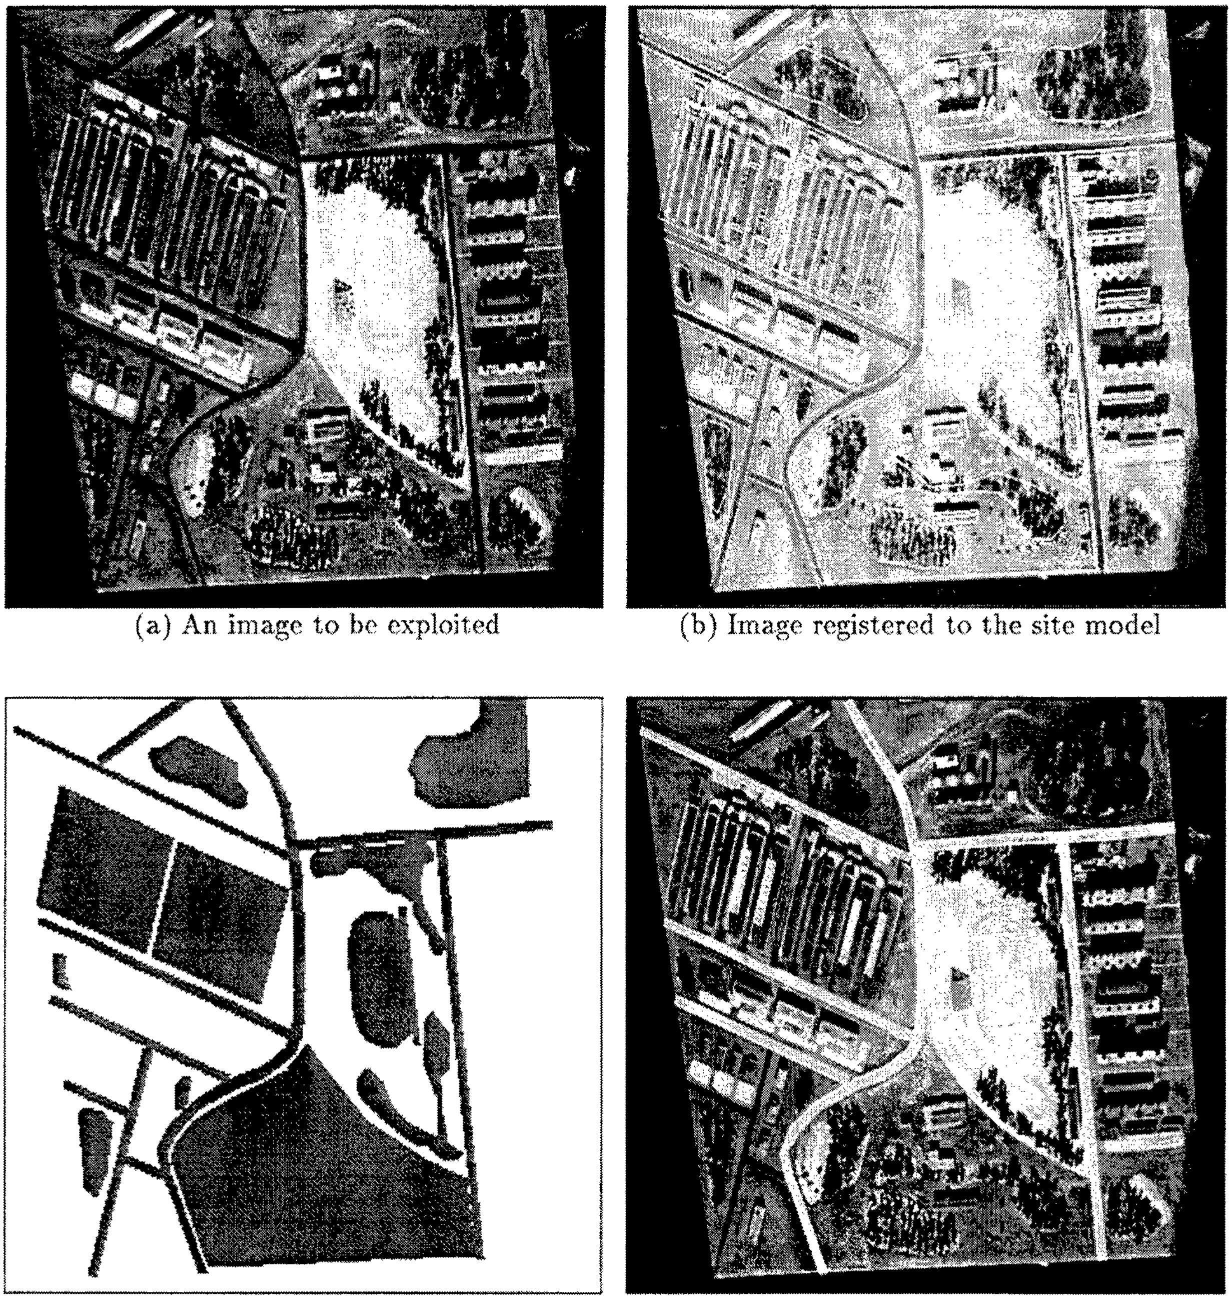

(a) An image to be exploited

(b) Image registered to the site model

(c) Regions of interest

(d) Regions to be monitored

Figure 9: Site model and ROIs.

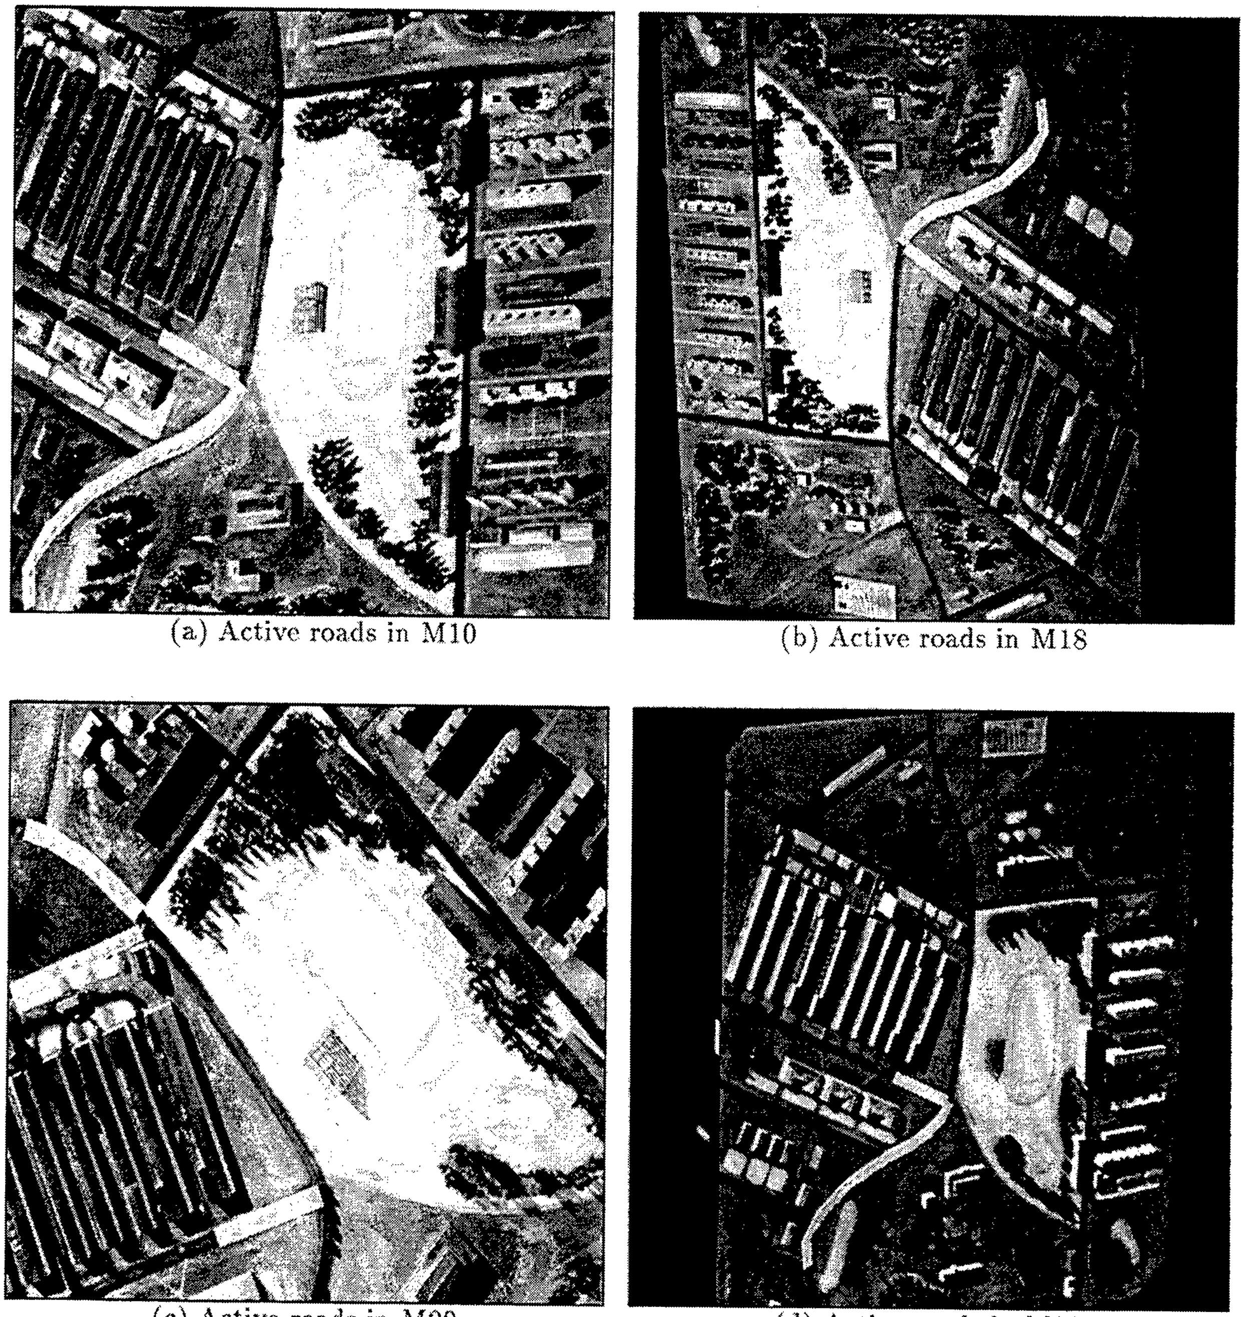

(a) Active roads in M10

(b) Active roads in M18

(c) Active roads in M20

(d) Active roads in M33

Figure 10: Convoy detection on M10, M18, M20, M33.

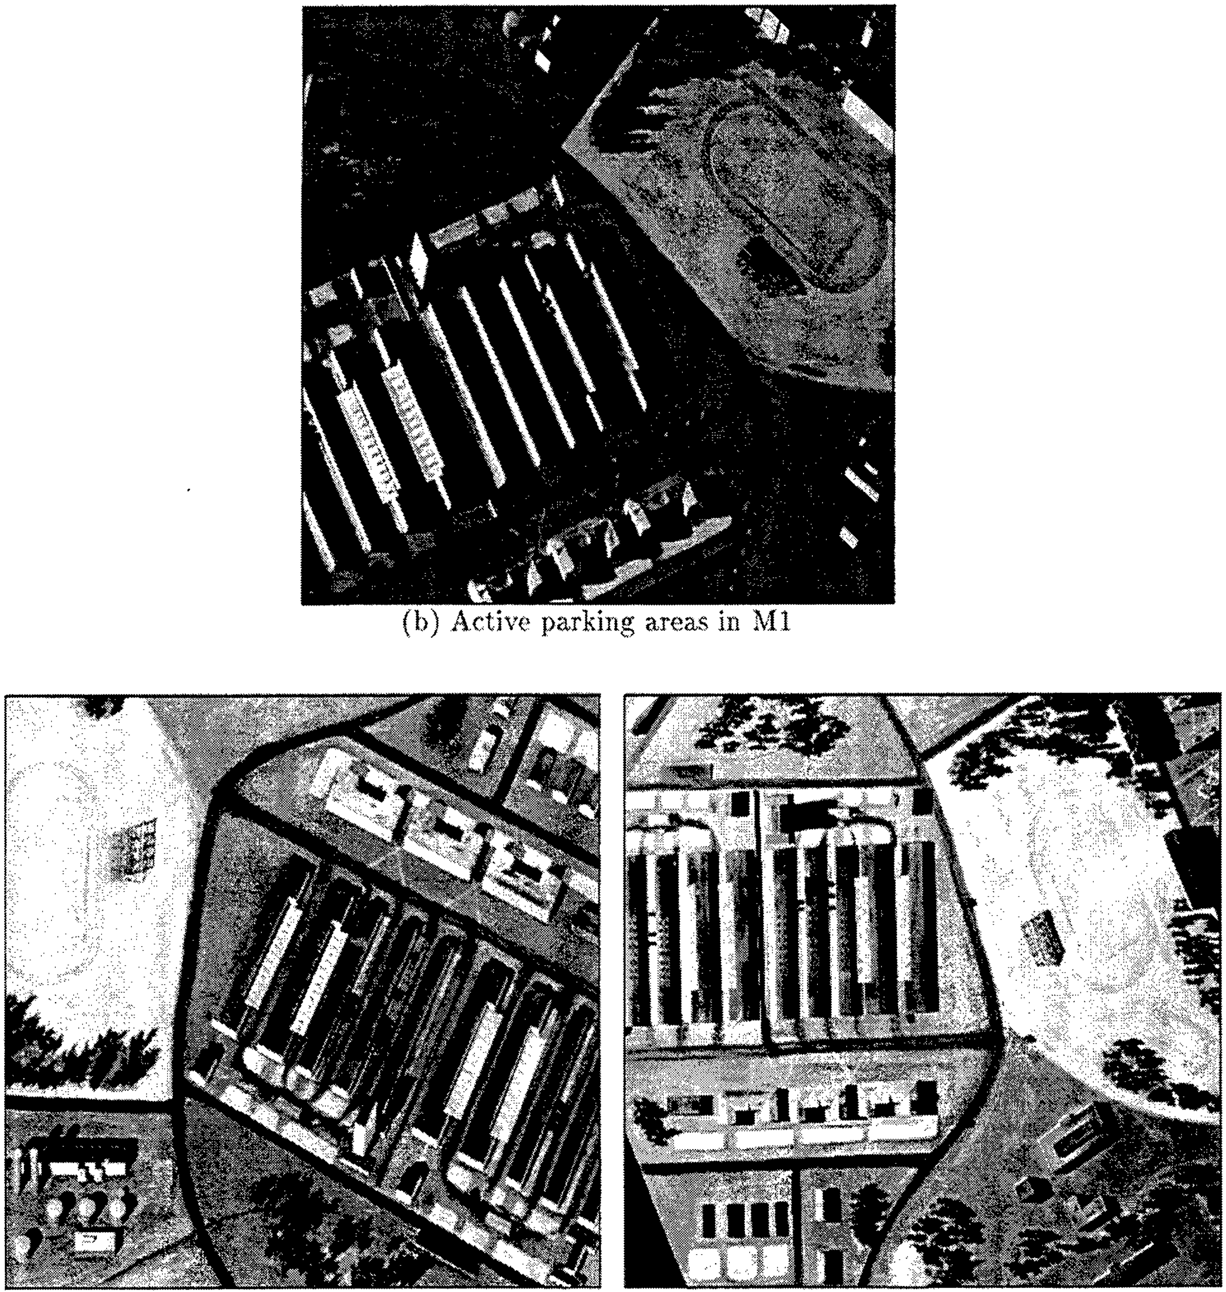

(b) Active parking areas in M1

(c) Active parking areas in M21

(d) Active parking areas in M11

Figure 11: Parking occupancy on M1, M21, M11.

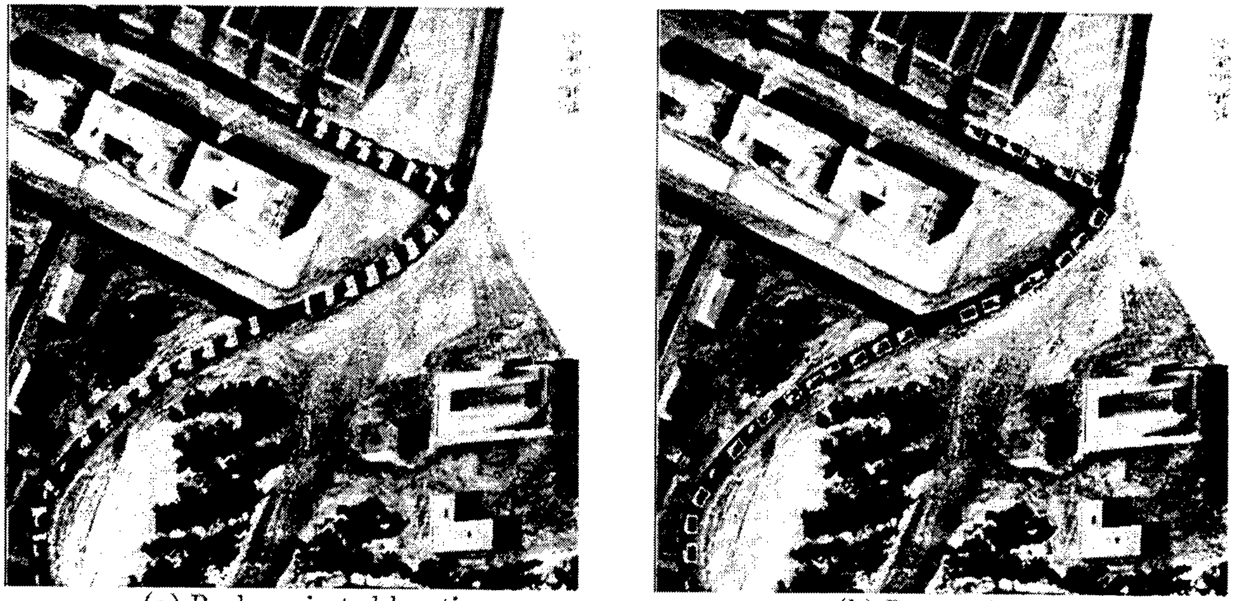

(a) Back-projected locations (b) Detected vehicles

Figure 12: Local vehicle detection.

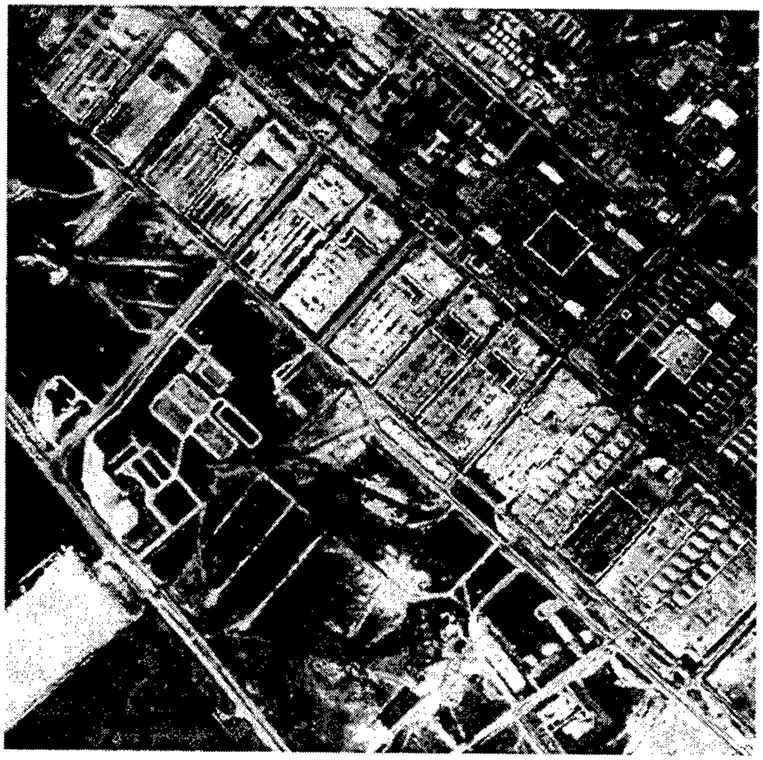

Figure 13: Ft. Hood image3.

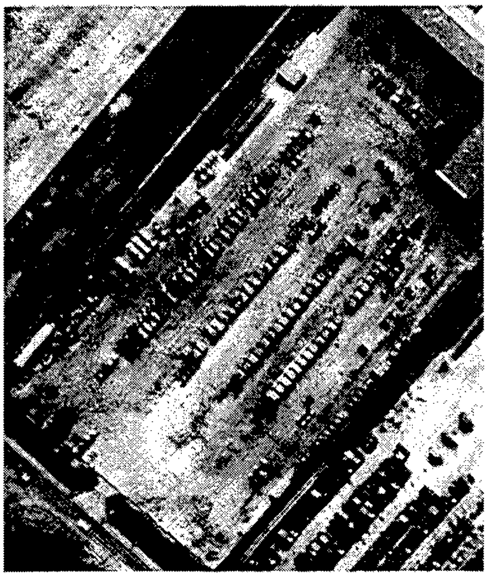

Figure 14: Region of interest.

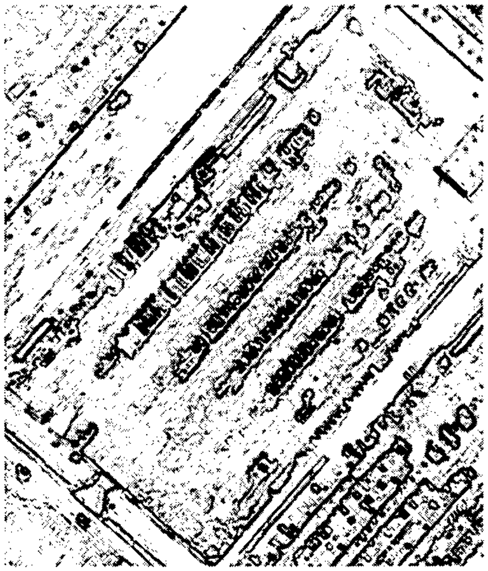

Figure 15: Edge map.

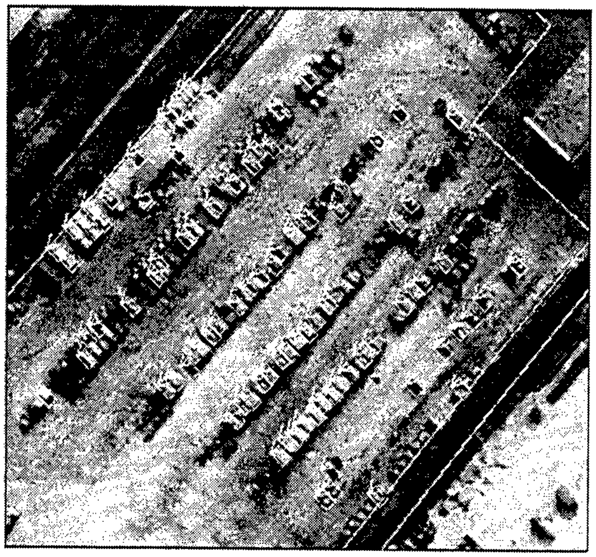

Figure 16: Detected vehicles.

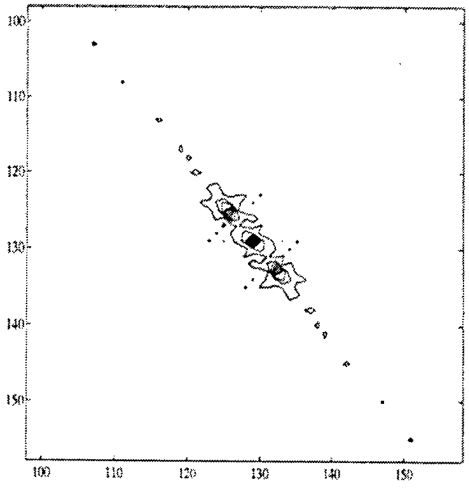

Figure 17: 2-D FT peak contours.

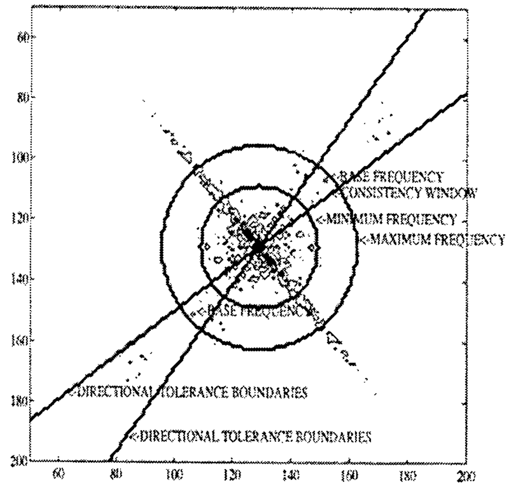

Figure 18: Search in compliance window.

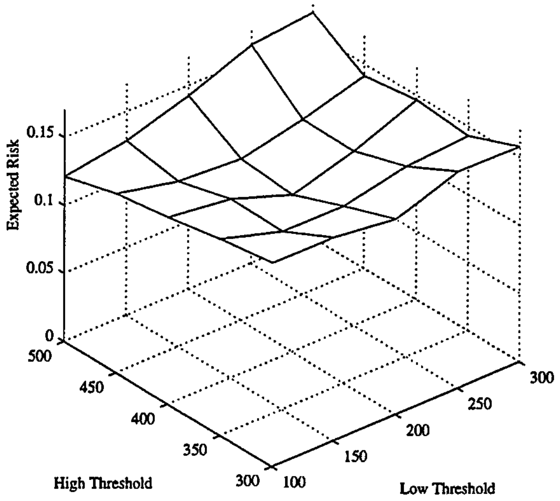

Figure 19: Vehicle detection: Expected risk as a function of Canny thresholds evaluated over the Ft. Hood and Denver training sets.

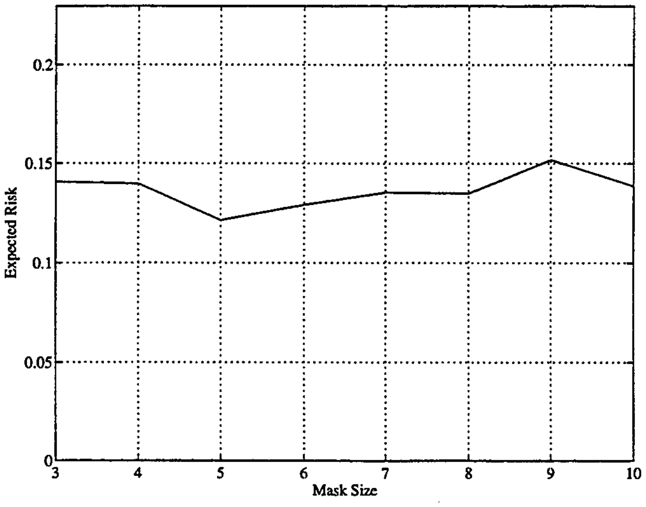

Figure 20: Vehicle detection: Expected risk as a function of Canny mask size evaluated over the Ft. Hood and Denver training sets.

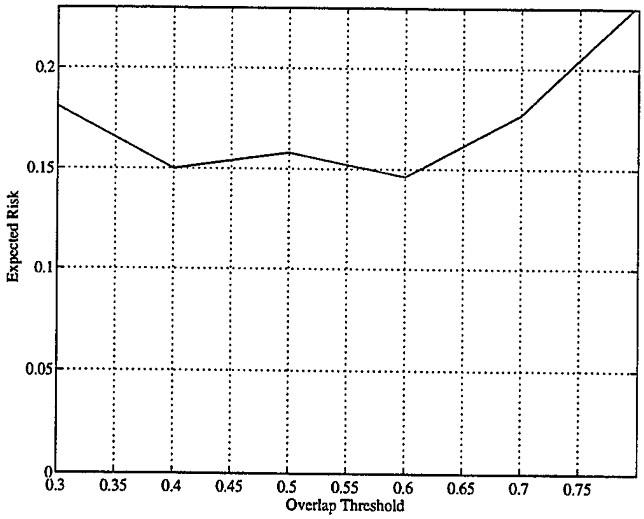

Figure 21: Vehicle detection: Expected risk as a function of overlap threshold evaluated over the Ft. Hood and Denver training sets.

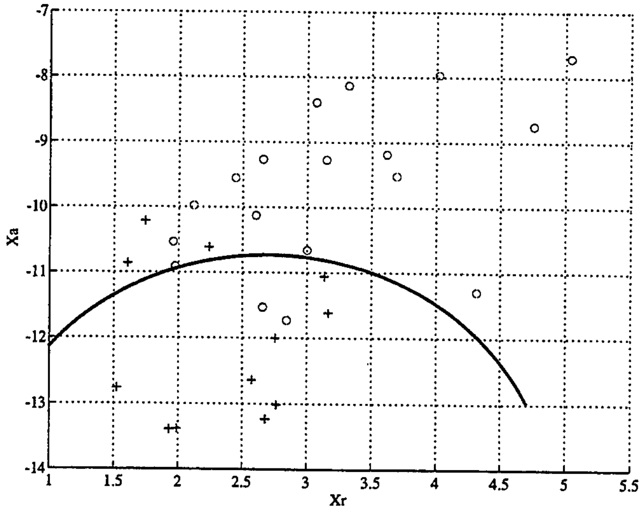

Figure 22: Decision region obtained from training images. Active parking lots are indicated by 'o's and inactive ones by '+'s in the (L_r, L_a) plane.

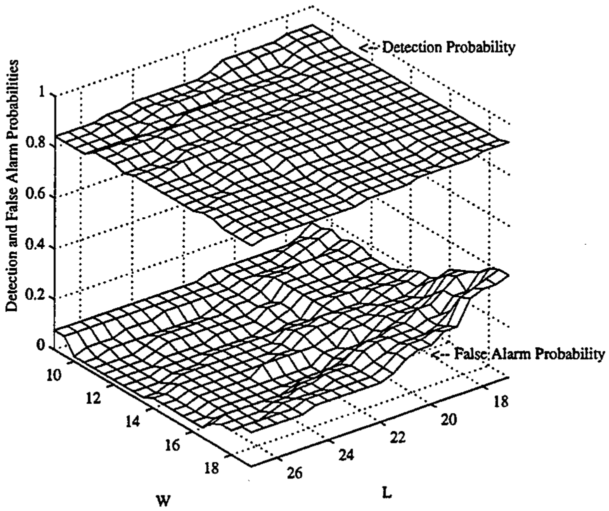

Figure 23: Sensitivity of the detection of active parking lots to misspecification of vehicle dimensions.

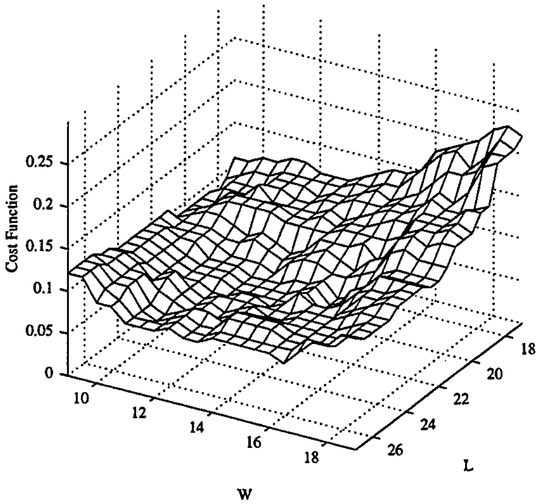

Figure 24: Risk function computed on all images for active parking lot detection.

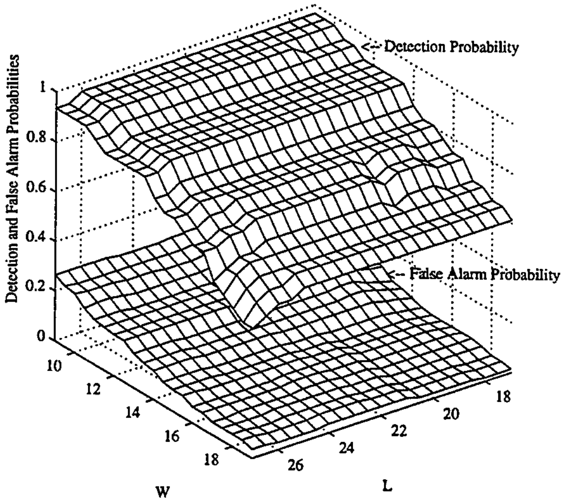

Figure 25: Sensitivity of convoy detection to misspecification of vehicle dimensions.

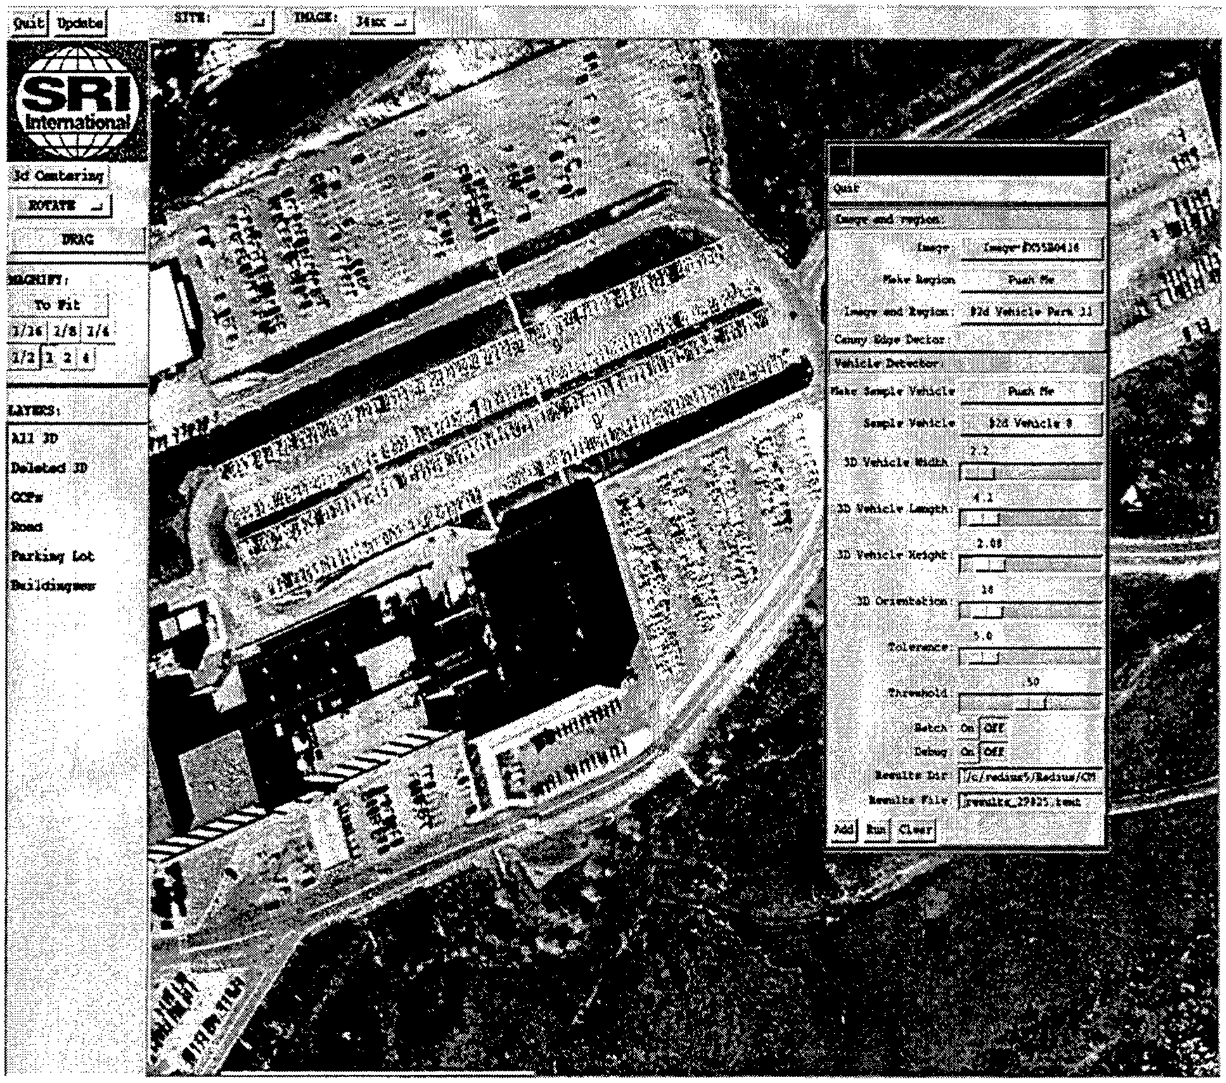

Image: Image-DC5B0418

Image and Region: $2d Vehicle Park 11

Sample Vehicle: $2d Vehicle 8

3D Vehicle Width: 2.2

3D Vehicle Length: 4.2

3D Vehicle Height: 2.08

3D Orientation: 18

Tolerance: 5.0

Threshold: .50

Batch: On [✓ OFF]

Debug: On [✓ OFF]

Results Dir: /c/redaux5/Radius/CM

Results File: results_29825.text

Figure 26: Integrated vehicle detector.

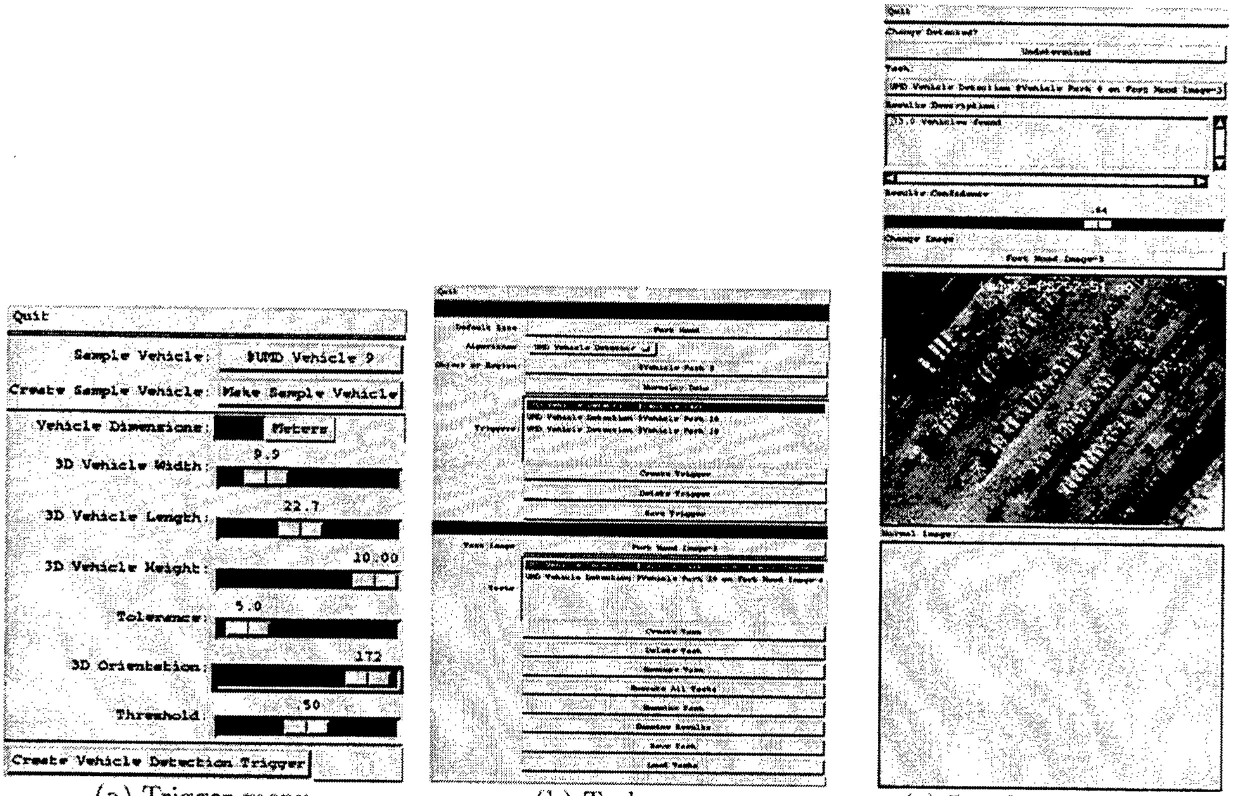

Sample Vehicle: $UMD Vehicle 9
Create Sample Vehicle: Make Sample Vehicle
Vehicle Dimensions: Meters
3D Vehicle Width: 9.9
3D Vehicle Length: 22.7
3D Vehicle Height: 10.00
Tolerance: 5.0
3D Orientation: 172
Threshold: .50

(a) Trigger menu

Default Site: Fort Hood
Algorithm: UMD Vehicle Detector
Object or Region: $Vehicle Park 9
Trigger: UMD Vehicle Detection $Vehicle Park 10
Task Image: Fort Hood Image-3
Task: UMD Vehicle Detection $Vehicle Park 10 on Fort Hood Image-4

(b) Task menu

Change Detected? Undetermined
Task: UMD Vehicle Detection $Vehicle Park 4 on Fort Hood Image-3
Results Description: 33.0 Vehicles found
Results Confidence: .64
Change Image: Fort Hood Image-3
Normal Image:

(c) Results window

Figure 27: Task and trigger, and result windows under the QL menus.

Table 1: Detection results on roads.

Image No.	Total Visible	Active Roads	Missed Roads	False Positive
3	8	0	0	1
4	6	1	0	0
5	7	1	1	0
7	5	0	0	0
8	11	0	0	3
10	10	3	0	0
11	9	0	0	0
13	13	0	0	0
16	11	4	0	5
17	11	0	0	1
18	12	4	1	0
20	7	3	0	0
21	9	0	0	0
26	9	3	0	3
27	12	0	0	0
32	12	5	1	0
37	9	2	1	0
40	10	5	1	2

Table 2: Detection of active parking areas: Hand-tuned rule.

Image No.	Total Visible	Active Areas	Missed Areas	False Positives	Image No.	Total Visible	Active Areas	Missed Areas	False Positives
1	4	2	0	0	21	4	4	0	0
2	3	1	0	0	22	4	4	0	0
3	2	2	0	0	23	2	0	0	1
4	2	0	0	0	24	2	2	0	0
5	2	0	0	0	25	2	2	2	0
6	2	0	0	0	26	3	1	0	0
7	2	0	0	0	27	4	4	1	0
8	2	2	0	0	28	2	0	0	0
9	2	2	1	0	29	4	2	0	1
10	3	1	0	0	30	4	2	0	1
11	4	4	0	0	31	4	4	0	0
12	4	4	0	0	32	2	1	0	1
13	3	3	0	0	33	4	2	0	0
14	4	4	0	0	34	3	1	0	2
15	3	1	0	0	35	4	4	0	0
16	4	2	0	0	36	2	0	0	0
17	4	2	0	1	37	3	1	0	0
18	4	2	0	2	38	2	0	0	0
19	4	2	0	0	39	4	2	1	0
20	2	0	0	0	40	3	0	0	0

Table 3: Sensitivity to Canny parameters: Detection performance vs. mask size.

		Mask size							
		3	4	5	6	7	8	9	10
TEC2-p1	Detection probability	0.78	0.79	0.79	0.77	0.74	0.73	0.72	0.74
	False positive rate	0.06	0.11	0.08	0.06	0.07	0.08	0.10	0.03
TEC2-p2	Detection probability	0.75	0.81	0.81	0.75	0.71	0.77	0.75	0.64
	False positive rate	0.05	0.01	0.03	0.07	0.07	0.01	0.03	0.00
TEC2-p3	Detection probability	0.88	0.90	0.90	0.92	0.88	0.80	0.82	0.90
	False positive rate	0.23	0.31	0.23	0.25	0.33	0.25	0.25	0.17
TEC2-p4	Detection probability	0.75	0.75	0.78	0.76	0.81	0.82	0.78	0.77
	False positive rate	0.10	0.10	0.10	0.11	0.08	0.06	0.06	0.06
TEC2-p5	Detection probability	0.78	0.73	0.76	0.75	0.76	0.76	0.76	0.72
	False positive rate	0.05	0.07	0.06	0.06	0.06	0.07	0.06	0.04
TEC2-p6	Detection probability	0.74	0.75	0.78	0.74	0.77	0.75	0.71	0.74
	False positive rate	0.08	0.06	0.04	0.03	0.04	0.06	0.08	0.06
Ft. Hood	Detection probability	0.90	0.92	0.92	0.95	0.95	0.90	0.90	0.92
	False positive rate	0.04	0.02	0.04	0.02	0.02	0.04	0.04	0.04

Table 4: Sensitivity to Canny parameters: Detection probability vs. threshold values.

			Low threshold				
			100	150	200	250	300
TEC2-p1 detection probability	High threshold	300	0.79	0.81	0.81	0.75	0.83
		350	0.79	0.81	0.79	0.79	0.82
		400	0.82	0.79	0.79	0.77	0.78
		450	0.79	0.81	0.81	0.81	0.77
		500	0.78	0.79	0.77	0.74	0.77
TEC2-p2 detection probability	High threshold	300	0.92	0.86	0.83	0.79	0.75
		350	0.84	0.86	0.81	0.79	0.73
		400	0.79	0.73	0.81	0.79	0.67
		450	0.77	0.71	0.69	0.69	0.60
		500	0.67	0.62	0.52	0.50	0.39
TEC2-p3 detection probability	High threshold	300	0.90	0.88	0.88	0.86	0.78
		350	0.90	0.90	0.92	0.88	0.78
		400	0.90	0.90	0.92	0.92	0.88
		450	0.84	0.94	0.90	0.92	0.82
		500	0.92	0.88	0.88	0.84	0.76
TEC2-p4 detection probability	High threshold	300	0.82	0.80	0.78	0.78	0.78
		350	0.81	0.78	0.79	0.77	0.77
		400	0.76	0.77	0.78	0.74	0.70
		450	0.76	0.77	0.74	0.70	0.69
		500	0.77	0.77	0.70	0.68	0.63
TEC2-p5 detection probability	High threshold	300	0.80	0.78	0.78	0.78	0.79
		350	0.76	0.76	0.78	0.75	0.77
		400	0.77	0.78	0.80	0.78	0.76
		450	0.80	0.81	0.81	0.80	0.77
		500	0.79	0.80	0.79	0.78	0.78
TEC2-p6 detection probability	High threshold	300	0.78	0.76	0.73	0.69	0.70
		350	0.71	0.74	0.72	0.72	0.72
		400	0.75	0.73	0.72	0.70	0.70
		450	0.76	0.72	0.75	0.70	0.69
		500	0.75	0.70	0.67	0.64	0.63
Ft. Hood detection probability	High threshold	300	0.92	0.92	0.92	0.90	0.90
		350	0.95	0.95	0.92	0.90	0.88
		400	0.92	0.97	0.95	0.92	0.92
		450	0.95	0.95	0.95	0.95	0.95
		500	0.95	0.97	0.97	0.95	0.95

Table 5: Sensitivity to Canny parameters: False positive rate vs. threshold values.

			Low threshold				
			100	150	200	250	300
TEC2-p1 false positive rate	High threshold	300	0.16	0.16	0.15	0.16	0.13
		350	0.15	0.15	0.13	0.12	0.10
		400	0.10	0.10	0.08	0.10	0.08
		450	0.13	0.10	0.07	0.06	0.05
		500	0.08	0.08	0.07	0.10	0.08
TEC2-p2 false positive rate	High threshold	300	0.30	0.26	0.18	0.15	0.15
		350	0.16	0.13	0.07	0.05	0.09
		400	0.07	0.11	0.01	0.00	0.05
		450	0.05	0.05	0.05	0.05	0.05
		500	0.01	0.03	0.07	0.03	0.00
TEC2-p3 false positive rate	High threshold	300	0.49	0.41	0.39	0.37	0.35
		350	0.37	0.33	0.27	0.29	0.25
		400	0.35	0.27	0.25	0.25	0.25
		450	0.27	0.27	0.19	0.17	0.11
		500	0.13	0.19	0.13	0.13	0.09
TEC2-p4 false positive rate	High threshold	300	0.06	0.08	0.06	0.07	0.08
		350	0.04	0.04	0.05	0.05	0.04
		400	0.08	0.08	0.08	0.06	0.10
		450	0.05	0.04	0.04	0.07	0.06
		500	0.07	0.06	0.07	0.08	0.07
TEC2-p5 false positive rate	High threshold	300	0.09	0.09	0.10	0.10	0.10
		350	0.08	0.08	0.08	0.09	0.07
		400	0.07	0.07	0.04	0.06	0.06
		450	0.06	0.04	0.07	0.07	0.08
		500	0.09	0.08	0.08	0.07	0.04
TEC2-p6 false positive rate	High threshold	300	0.04	0.06	0.03	0.04	0.05
		350	0.08	0.04	0.04	0.04	0.03
		400	0.03	0.03	0.03	0.02	0.04
		450	0.03	0.04	0.03	0.01	0.01
		500	0.04	0.03	0.02	0.03	0.03
Ft. Hood false positive rate	High threshold	300	0.04	0.04	0.02	0.04	0.04
		350	0.04	0.02	0.02	0.02	0.02
		400	0.07	0.02	0.02	0.04	0.04
		450	0.02	0.00	0.00	0.00	0.00
		500	0.00	0.00	0.00	0.04	0.04

Table 6: Sensitivity to overlap threshold: Detection performance vs. threshold.

		Threshold					
		0.3	0.4	0.5	0.6	0.7	0.8
TEC2-p1	Detection probability	0.84	0.83	0.82	0.81	0.75	0.54
	False positive rate	0.29	0.21	0.10	0.04	0.01	0.02
TEC2-p2	Detection probability	0.86	0.81	0.71	0.71	0.62	0.50
	False positive rate	0.08	0.04	0.07	0.05	0.02	0.00
TEC2-p3	Detection probability	0.90	0.90	0.86	0.86	0.82	0.76
	False positive rate	0.34	0.24	0.22	0.18	0.12	0.02
TEC2-p4	Detection probability	0.81	0.81	0.74	0.73	0.69	0.61
	False positive rate	0.18	0.09	0.11	0.09	0.08	0.06
TEC2-p5	Detection probability	0.78	0.76	0.74	0.72	0.70	0.60
	False positive rate	0.12	0.11	0.09	0.07	0.07	0.06
TEC2-p6	Detection probability	0.78	0.82	0.72	0.72	0.64	0.54
	False positive rate	0.13	0.06	0.04	0.02	0.02	0.02
Ft. Hood	Detection probability	0.88	0.88	0.90	0.90	0.69	0.45
	False positive rate	0.26	0.15	0.07	0.05	0.06	0.09

Table 7: Detection of active parking areas: Learned rule.

Image No.	Total Visible	Active Areas	Missed Areas	False Positive	Image No.	Total Visible	Active Areas	Missed Areas	False Positive
1	4	2	0	0	21	4	4	0	0
2	3	1	0	0	22	4	4	0	0
3	2	2	0	0	23	2	0	0	0
4	2	0	0	0	24	2	2	0	0
5	2	0	0	0	25	2	2	2	0
6	2	0	0	0	26	3	1	0	0
7	2	0	0	0	27	4	4	0	0
8	2	2	1	0	28	2	0	0	0
9	2	2	0	0	29	4	2	0	0
10	3	1	0	0	30	4	2	0	0
11	4	4	0	0	31	4	4	0	0
12	4	4	0	0	32	2	1	0	0
13	3	3	0	0	33	4	2	0	0
14	4	4	0	0	34	3	1	0	0
15	3	1	0	0	35	4	4	0	0
16	4	2	0	0	36	2	0	0	0
17	4	2	0	1	37	3	1	0	0
18	4	2	0	2	38	2	0	0	0
19	4	2	2	0	39	4	2	1	0
20	2	0	0	0	40	3	0	0	0